Moab Cluster Manager User's Guide

Cluster Resources

Moab Cluster Manager User's Guide
by Cluster Resources

Copyright © 1999 - 2007 Cluster Resources, Inc.

Table of Contents

Notice ..... **vii**

1. Getting Started **1**

1.1. Getting Started Overview 1
1.2. Installation 1
1.2.1. Unix-Based installation 1
1.2.2. Windows installation 1
1.3. Connection Wizard 2
1.3.1. Remote Connection 2
1.3.2. Local Connection 3
1.3.3. Offline Demonstration 4
1.3.4. Online Demonstration 4
1.4. View Summary 4
1.4.1. Main Menu Bar 4
1.4.2. Dashboard 5
1.4.3. Main Info Screen 5
1.4.4. System Utilization Bar 7

2. Workload **8**

2.1. Workload Overview 8
2.2. Jobs 8
2.2.1. Create Job 8
2.2.2. Modify A Job 15
2.2.3. List Jobs/Job Templates 26
2.2.4. Job Timeline 39
2.2.5. Job Outlines 39
2.2.6. Dynamic Job Allocation 40
2.3. Reservations 40
2.3.1. Create Reservation 40
2.3.2. Modify Reservation 47
2.3.3. List Reservations 51
2.3.4. List Recurring Reservations 56
2.3.5. Reservation Timeline 57
2.3.6. Reservation Calendar 57
2.4. Triggers 58
2.4.1. List Triggers 58

3. Resources **61**

3.1. Resources Overview 61
3.2. Moab Workload Manager 61
3.2.1. Control Panel 61
3.2.2. Log Settings 61
3.2.3. System Settings 64
3.2.4. Simulation Settings 66
3.2.5. Statistics Settings 68
3.2.6. High Availability 68
3.3. Nodes 69
3.3.1. Create Node Profile 69

3.3.2. Modify A Node Profile 71
3.3.3. List Nodes 72
3.3.4. Node Calendar 77
3.3.5. Node Timeline 78
3.4. Partitions 78
3.4.1. Create a Partition Profile 78
3.4.2. Modify a Partition Profile 79
3.4.3. List Partitions 79
3.5. Licenses 81
3.5.1. List Licenses 82
3.6. Cluster 83
3.6.1. Visual Cluster 83
3.6.2. Processor Usage 88
3.6.3. Add Resource Manager 88
3.6.4. List / Modify Resource Managers 89
3.6.5. Resource Manager Messages 90
3.6.6. Allocation Manager 91
3.7. Grid 93
3.7.1. Grid Summary 94
3.7.2. Visual Grid 96
3.7.3. Create Grid Relationship 97
3.7.4. Modify Grid Relationship 99
4. Organization 102
4.1. Organization Overview 102
4.2. Visual Credential Access 102
4.3. User Access 102
4.4. List Credentials 103
4.5. Create/Modify a User Profile 109
4.6. Create/Modify a Group Profile 114
4.7. Create/Modify an Account Profile 118
4.8. Create/Modify a Class Profile 123
4.9. Create/Modify a Quality of service Profile 127
5. Policies 136
5.1. Policies Overview 136
5.2. Fairshare 136
5.2.1. Fairshare Options 136
5.2.2. Fairshare Table 137
5.3. Priority 137
5.3.1. Main Priority Components 138
5.3.2. Wait Time Job Services 139
5.3.3. QoS Targets 139
5.3.4. Credential Priority 140
5.3.5. Job State Priority 140
5.3.6. Fairshare Usage Priority 140
5.3.7. Resource Requests Priority 141
5.3.8. Executing Job Usage Priority 141
5.3.9. Priority Display Options 141

5.4. Job Policies and Settings 142
5.4.1. Job Limit Settings 142
5.4.2. Job Defer Settings 142
5.4.3. Global Job Policy Settings 143
5.5. Reservation Policies 143
5.6. Resource Violation 145
5.7. Node Policies 147
5.8. Partition Policies 151
5.8.1. Partition Allocation Policy 151
5.9. Backfill 151
5.10. Role Based Authorization 154
6. Statistics 157
6.1. Statistics Overview 157
6.2. Quick Charts/Graphs 157
6.3. Matrix Statistics 157
6.4. Custom Charts/Graphs 159
6.4.1. Credential Based Charts 159
6.4.2. Node Categorization Charts 161
6.4.3. Generic Metric Charts 161
6.4.4. Job Template Charts 162
6.5. Custom Reports 162
7. Diagnostics 164
7.1. Diagnostics Overview 164
7.2. Diagnostics Support 164
8. Miscellaneous 165
8.1. Miscellaneous Overview 165
8.2. Console 165
8.3. Debugging and Log Levels 165
8.4. Cluster Manager Preferences 166
8.5. Plugin Manager 166
9. License 168

List of Tables

3-1. Visual Cluster Example 83

Notice

Important: This is the general release of the Moab Cluster Manager User's Guide. Other information may be found by browsing the Cluster Resources website at http://www.clusterresources.com/

Chapter 1. Getting Started

1.1. Getting Started Overview

Moab Cluster Manager (MCM) is a Java-based graphical interface for managing the Moab Workload Manager. It allows users to submit jobs, schedule reservations, view job statistics, etc. in an easy, user-friendly way.

This chapter explains how to get started using the Moab Cluster Manager by installing it, connecting it to a Moab Workload Manager, and describing it's main window.

1.2. Installation

1.2.1. Unix-Based installation

Moab Cluster Manager is written in Java 1.5. You will need to have Java 1.5 or higher installed on your system to run MCM. This may be an existing copy on your system, or the JRE bundled with the MCM distribution. These install instructions assume a basic familiarity with Unix/Linux file systems, and commands such as ls, tar, mv, etc.

The installation steps are as follows:

1. Download the a version of the tar file from the Cluster Resources web site:
   a. The mcm-[version]-[build number].linux.tar.gz comes with a bundled JRE.
   b. The mcm-[version]-[build number].tar.gz does NOT include a JRE.
2. Move the tar file to your home directory, or another directory you have access to (i.e. /home/username/). Unpack the tar file (tar -xzvf xxxx.tar.gz).
3. Change to the newly unpacked MCM directory.
4. You may now start MCM at anytime by running the mcm script (i.e. ./mcm). This script will check for the existence of Java and then run MCM.

1.2.2. Windows installation

1. Download the installation executable from the Cluster Resources web site. The mcm-[version]-[build number].exe is the Windows installer that will setup MCM on your system.
2. Double-click the installation file. The MCM installer will guide you through the installation process.

Note: The default target folder is "C:\Program Files\Moab Cluster Manager".

3. The installer will create Start Menu and Desktop icons that can be used to run MCM.
4. Double-click the Moab Cluster Manager icon on the Desktop to run MCM.

1.3. Connection Wizard

The Connection Wizard provides four connection options for the Moab Cluster Manager: Remote Connection, Local Connection, Offline Demonstration and Online Demonstration.

1. Remote Connection: Connects to a remote Moab Workload Manager server over SSH.
2. Local Connection: Connects to a locally hosted Moab Workload Manager server.
3. Offline Demonstration: Allows a user to view a demonstration snapshot.
4. Online Demonstration: Automatically logs in to an online demonstration cluster for a preview of the product.

1.3.1. Remote Connection

The Remote Connection feature allows you to securely connect to a remote Moab Workload Manager server. Here is a description of each of the connection options:

- Host Name & Port
  - Host Name (or IP Address) - The Host Name (or IP address) of the server that is running Moab Workload Manager. If you do not know the host name or IP address of the server, please consult your system administrator.
  - Port - The port on which SSH is running on the remote server (the default is 22). If you do not know which port to use, please consult your system administrator.

- Authentication Options
  - Password Authentication - This option tells Moab Cluster Manager to authenticate by prompting the user for a password. Consult with your system administrator for information regarding your user name, password and the type of authentication used.
  - Keyboard Interactive Authentication - This option tells Moab Cluster Manager to interactively prompt for authentication information.
  - SSH Key Authentication - This option tells Moab Cluster Manager to connect to the remote computer using only the user name and a private key file. Consult your system administrator for information regarding your user name, private key and the type of authentication used.

  - Ask for SSH Key Passphrase - Some private keys require a passphrase to be entered before it will allow a user to authenticate. In this is the case this box should be checked, otherwise an empty passphrase will be used for authentication.

- Connection Settings
  - User Name - This is the name used to login to the remote computer. Consult your system administrator for information regarding your user name or password.
  - Path to Moab Workload Manager Client Commands (ie showq) - The directory containing the Moab Client Commands, such as showq, mschedctl, mdiag, etc. This is not the location of the Workload Manager, but instead the location of the commands that control the Workload Manager. This location is usually "/usr/local/bin".
  - Private Key Path - If SSH key Authentication is being used, this field is for the path of the private key file.

- Load, save or delete stored sessions
  - Auto Connect On Next Session - This option sets the Moab Cluster Manager to automatically connect to the specified saved session the next time it is run.
  - Saved Sessions - This field is where a user is able to save his/her remote connection settings so that they don't need to be entered each time.
  - Load Button - This button will load the selected saved session in the list to the left.
  - Save Button - This button will save a session according to the name typed in the Saved Sessions field.
  - Delete Button - This button will delete the saved session that is selected in the list to the left.

- Open Button - This button will open a connection to a remote Moab Workload Manager. It will attempt to authenticate the user based on the given settings, and then open the Moab Cluster Manager.
- Cancel Button - Clicking this button will close the Moab Cluster Manager Connection Wizard

1.3.2. Local Connection

This option connects to a Moab Workload Manager running on the local machine. The only option for this mode is the Path for the Moab Workload Manager Client Commands.

- Path to Moab Workload Manager Client Commands (ie showq) - The directory containing the Moab Client Commands, such as showq, mschedctl, mdiag, etc. This is not the location of the Moab Workload Manager, but instead the location of the commands that interact with the Moab Workload Manager. This location is usually "/usr/local/bin".
- Open Button - This button will open a connection to a local Moab Workload Manager and then open the Moab Cluster Manager.
- Cancel Button - Clicking this button will close the Moab Cluster Manager Connection Wizard

1.3.3. Offline Demonstration

Moab Cluster Manager is capable of recording all the data gathered from a cluster and saving it to a demonstration snapshot. This connection option will allow a user to view a previously recorded demonstration snapshot. Here is a description of what each of the buttons on this screen do:

- Import Button - The import button allows the user to select a file to copy to the appropriate Moab Cluster Manager directory. This allows users to import demonstration snapshots from other Moab Cluster Managers.
- Delete Button - This button will delete a saved demonstration snapshot.
- Open Button - This button will open the selected Moab Cluster Manager demonstration snapshot.
- Cancel Button - Clicking this button will close the Moab Cluster Manager Connection Wizard

1.3.4. Online Demonstration

The Online Demonstration is a free, online demonstration cluster for users to preview. This connection option will automatically log in to the demonstration cluster.

- Open Button - This button will connect Moab Cluster Manager to the Cluster Resources demonstration Moab Workload Manager.
- Cancel Button - Clicking this button will close the Moab Cluster Manager Connection Wizard

1.4. View Summary

The Moab Cluster Manager main window provides an overview of the current state of the cluster. There are four parts of this main window: the Main Menu Bar, the Dashboard, the Main Info Screen, and the System Utilization Bar.

1.4.1. Main Menu Bar

The Main Menu Bar is located across the top of the Moab cluster Manager window. Through this menu (File, Configure, Manage, etc.), all Moab Cluster Manager features can be accessed. The services unique to this menu are:

- Console window
- Save System Snapshot window
- Moab Cluster Manager Prefrences window
- Plugin Manager window

- About window

1.4.2. Dashboard

The Dashboard is a directory of all the services that the Moab Cluster Manager can provide to users and administrators. The availability of some services depends on the user's privileges, as determined by the ADMINCFG level defined in the moab.cfg file. The chapters of this User Guide mimic the layout found in the Dashboard.

1.4.3. Main Info Screen

The largest area in the Moab Cluster Manager main window is the Main Info Screen. This screen is intended to give general information about the system that Cluster Manager is currently connected to.

1.4.3.1. Scheduler Information

This panel displays the following information about the scheduler:

- **Name -** The name of the scheduler. Has no impact on operation.
- **Host -** This refers to the host computer where the Resource Manager is running.
- **Port -** The specific port that the scheduler is operating on.
- **Mode -** The operating mode of the scheduler. Mode options are shown in this table:
  -

NORMAL (default)	Normal operation, controls the Resource Manager as configured
MONITOR	MONITOR mode behaves identical to NORMAL mode except the ability to start, cancel, or modify jobs is disabled. This allows safe diagnosis of the scheduling state and behavior using the various diagnostic client commands.
INTERACTIVE	Like NORMAL mode, except Moab sends the desired change request to the screen and asks for permission to complete it.
SIMULATION	Processes a simulated environment as specified in the Workload.Trace and Resource.Trace files.

- **Status -** Indicates whether the scheduler is running, down, or paused.

1.4.3.2. Node Summary

This panel displays a high-level view of the state of the nodes found within the cluster. Click on any label to obtain a detailed list of nodes in the given cateogry. Category Descriptions:

- **Busy Nodes -** *Busy* nodes include all nodes which are actively executing batch jobs. A node will be listed as busy even if it is only partially loaded with jobs.
- **Idle Nodes -** *Idle* nodes include all nodes which are available but are currently not running any jobs.
- **Down Nodes -** *Down* nodes include all nodes which have reported major software, hardware, or batch failures or have been marked 'down' or 'offline' by an administrator.
- **Total Nodes -**The *total* nodes category includes all nodes in the cluster and is a sum of the *busy*, *idle*, and *down* nodes listed above.

1.4.3.3. Job Summary

This panel displays a high-level view of the state of the jobs found within the cluster. Click on any label to obtain a detailed list of jobs in the given cateogry. Category Descriptions:

- **Running Jobs -** *Running* jobs include all jobs which are actively executing or performing post-execution clean-up. This includes jobs in the states *starting*, *running*, or *exiting*.
- **Eligible Jobs -** *Eligible* jobs include all jobs which are in state *idle* and are not blocked by holds, usage limits, or other policies. Eligible jobs typically will run as soon as resources become available and may already have a reservation in place.
- **Blocked Jobs -** *Blocked* jobs include all jobs which cannot run due to reasons other than resource availability. Jobs may be blocked by job holds, resource manager level policies, scheduler job usage policies, job deadline constraints, or other factors. Clicking on the label will bring up the detailed blocked job list which will include additional information in the *blocked reason*column.
- **Total Jobs -** The *total* jobs category includes all jobs in the cluster and is a sum of the *running*, *eligible*, and *blocked*jobs listed above.

1.4.3.4. User Information

This panel displays the user information of whoever started the Moab Cluster Manager.

- **User -** The name of the user running Moab Cluster Manager.
- **Group -** The name of the user's Group.
- **Account -** Any accounts the user may belong to.
- **Class -** Classes the user has access to.

- **QoS -** Any QoS accounts the user may belong to.

1.4.3.5. User Job Summary

Displays information concerning jobs run by the current user.

- **Running Jobs -** Jobs run by the current user.
- **Eligible Jobs -** Users jobs that are waiting in the queue.
- **Blocked Jobs -** Users jobs that have been blocked, either by policy or user.
- **Total Jobs -** Total jobs from the user.

1.4.4. System Utilization Bar

The System Utilization bar displays *historical* system utilization as it pertains to utilized processors as captured by Moab profiling intervals. Also note that processors utilized are only measured once a job is finished.

Chapter 2. Workload

2.1. Workload Overview

The workload category of features deals with the submitting and viewing of jobs, reservation and triggers. These functions are used to get work done by the system.

2.2. Jobs

2.2.1. Create Job

A cluster runs programs. A job tells a cluster when, where, and how to run the programs. The create job window, often referred to as a job submission window, is how a user creates a job.

2.2.1.1. Job Creation

Job Information		
Field	**Required/Optional**	**Field Information**
Script / Executable / Program / Application	Required	A job consists of a script, executable, program or application. In order for the job to start it is necessary for it to know the location of the program. This field allows the user the ability to specify that location. Consult your system administrator for more specific information regarding your program's location.
Program Arguments / Options	Optional	Some programs provide users with different options. This field allows the user the ability to specify those options. A user should consult his/her program documentation to learn about the available options.

Job Information		
Job Name	Optional	A user can attach a custom name to a job to assist him/her in identifying the job. The name is provided only for the user's convenience and does not affect any policies or settings.
Template(s)	Optional	If there are job templates that are selectable, the user can do so here. Any attributes associated with the job templates will be mapped onto the submitted job.
User Job Priority	Optional	The higher a job's priority the sooner it will start. By changing this field a user can reduce their job's priority and change the order in which their job starts. This field is usually utilized to execute a users jobs in a specific order. This field only supports negative numbers with the exception of 0. A 0 User Job Priority will not delay the job from starting. However, the higher a negative number the lower a job's priority. For example, a job with a User Job Priority of -100 will allow more jobs to start before it starts, thus postponing the job from starting, then a user job priority of -10.
Administrator Job Priority	Optional	If this field is changed from zero it becomes a special administrator job. An administrator jobs starts before all other jobs with the exception of other administrator jobs. Note: only users with admin1 rights can create an administrator job.

Resources		
Field	**Required/Optional**	**Field Information**

Resources		
Cluster/Partition	Optional	Clusters are often divided into different sections. These sections are commonly called partitions. In a grid, clusters are also considered partitions. Users can only request one specific partition for their job. Consult your system administrator to learn which partition is the best suited for your job.
Reservation	Optional	A user can specify a reservation for this job. If a reservation is specified the job will execute only on the nodes that are reserved by the reservation.
Grid Policy	Optional	Sometimes a user has a program/script/executable/application that requires information from another program/script/executable/application before it can start. This field allows a user to specify a job that must finish execution before this job will be eligible to start.
Job Dependency	Optional	Sometimes a user has a program/script/executable/application that requires information from another program/script/executable/application before it can start. This field allows a user to specify a job that must finish execution before this job will be eligible to start.
Node Features	Optional	Some jobs require a specific feature on a node. A feature is a custom tag attached to a specific list of nodes. Consult your system administrator for specific information regarding each tag.

Resources		
Node List	Optional	A node is a computer consisting of 1 or more processors. This field allow the user to define which nodes a job requires to execute. If a node list is not specified the nodes needed for the job are gathered from the nodes field.
Nodes	Optional	A node is a computer consisting of 1 or more processors. A job requires at least 1 processor to execute and therefore must use at least 1 node. If this field is set to 0 Workload Manager assigns the job to 1 node unless the *Node List*field is populated.
Processors	Optional	All jobs requires at least 1 processor. If this field is not used the processors are calculated by using the available processors on a node. If a processor and a node are not requested Workload Manager assigns 1 node to the job.
Memory per Node	Optional	Some jobs require specific amounts of memory. This field allows a job to request the memory it needs for each node. It should be noted that this field is not the total memory across the entire cluster but only the memory on each node needed by the job. Workload Manager will start this job only on the nodes that have sufficient memory. If this field is not used then Workload Manager will start the job on any available node.

Resources		
Swap per Node	Optional	Some jobs require specific amounts of swap. This field allows a job to request the swap it needs for each node. It should be noted that this field is not the total swap across the entire cluster but only the swap on each node needed by the job. Workload Manager will start this job only on the nodes that have sufficient swap. If this field is not used then Workload Manager will start the job on any available node.
Operating System	Optional	If an operating system is selected, Moab will try to run the job on any nodes with the specified operating system.
Architecture	Optional	If an architecture is selected, Moab will try to run the job on any nodes with the specified architecture.
Search Resources	Optional	This button displays a table allowing the user the ability to search for available resources. Consult the Search Resources documentation for more specific information.
Estimated Start Time Calculator	Optional	This button displays Moab's text based output determining when a job can start.
Estimated Start Time Table	Optional	This button displays a table of the estimated start times for jobs of different processor sizes. Consult the Estimated Start Time documentation for more specific information.

Time Frame		
Field	**Required/Optional**	**Field Information**

Time Frame		
Start Time	Optional	Some jobs require a specific amount of time before they should be allowed to start. This field allows the user the ability to define when the earliest time that the job can begin. By default, a job may start as soon as resources become available.
Duration	Optional	The duration is the estimated time needed for a job to execute. If a user's job requires more time than the specified duration, duration violation policies come in to effect. Consult your system administrator for more information regarding these policies. If no duration is specified, a default wall time will be applied. Consult your system administrator for more information regarding your cluster's default wall time.

Job Environment		
Field	**Required/Optional**	**Field Information**
Execution Directory	Optional	Some jobs need to be executed in a specific location on each node. This field allows a user the ability to define that location. By default, the job is executed in the user's home directory. Consult your system administrator for information regarding your home directory.
Input File	Optional	Some scripts/executables/programs/applications required input files to be able to execute. This field allows the user the ability to define those files.

Job Environment		
Output Directory	Optional	All scripts/executables/programs/applications use an output directory. This field gives the user the ability to customize the location of the output directory.
Error Directory	Optional	All scripts/executables/programs/applications use an error directory. This field gives the user the ability to customize the location of the error directory.

Credential Information		
Field	**Required/Optional**	**Field Information**
User	Required	This field defines the name of the user under whom this job will execute. Only users with Admin1 rights can change this field.
Group	Optional	This field defines the name of the group under whom this job will execute. Only the groups available to the user are displayed.
Account	Optional	This field defines the name of the account under whom this job will execute. Only the accounts available to the user are displayed.
Class	Optional	This field defines the name of the class under which this job will execute. Only the classes available to the user are displayed.
Quality of Service (QoS)	Optional	This field defines the name of the quality of service (QoS) under which this job will execute. Only the QoS's available to the user are displayed.

Job Flags		
Field	**Required/Optional**	**Field Information**

Job Flags		
Hold	Optional	A hold can only be placed upon jobs that haven't began execution. A hold stops or halts a job from running until the user or an administrator releases the hold.
Preemptible	Optional	A job that is preemptible can be suspended or re-queued by higher priority jobs.
Preemptor	Optional	The job may preempt other jobs which have the **PREEMPTEE** flag.
Restartable	Optional	If a job experiences a failure during execution the user must resubmit the job to Workload Manager. However, a job that is restartable would automatically be restarted by Workload Manager in the event of a failure.

Email Notification		
Field	**Required/Optional**	**Field Information**
Job Completion	Optional	When a job finishes execution an email notification will be sent to the user stating this. **NOTE**: User email addresses may be specified in the 'Create/Manage Users' page.
Job Start	Optional	When a job begins execution an email notification will be sent to the user stating this. **NOTE**: User email addresses may be specified in the 'Create/Manage Users' page.
Job Failure	Optional	When a job cannot start or crashes during execution an email notification will be sent to the user stating this. **NOTE**: User email addresses may be specified in the 'Create/Manage Users' page.

2.2.2. Modify A Job

Summary:

A cluster runs programs. A job tells a cluster when, where, and how to run the programs. The modify job window allows a user to modify an already existing job. The fields that can be modified are user job priority, system priority, and duration. If multiple jobs are selected, then QoS can also be modified.

> *Information*		
Field	**Displayed**	**Field Information**
Job ID	Always	All jobs, when created, are given a unique ID by Workload Manager. This field displays that ID.
Job Name	Data Dependent	Users can attach a custom name to the job to allow them to easily identify their jobs. The name does not change any Workload Manager settings or prioritizations. If a name has been attached, it will appear in this field.
Hold	Data Dependent	A hold can only be placed upon jobs that haven't started. A hold stops or halts a job from running until the user or an administrator releases the hold. If a hold has been placed it will be displayed in this field.
State	Always	This field will display the execution status of the job. For example, running, stopped, executing, idle, blocked, etc.
Messages	Data Dependent	This field will display informational messages relating to the job.
Blocked Reason / Error	Data Dependent	This field will display diagnostic messages relating to the job.

> *Credentials*		
Field	**Displayed**	**Field Information**
User	Always	This field will display the user ID under which the job is executing.

> *Credentials*		
User Job Priority	Data Dependent	The higher a job's priority, the sooner it will start. A user has the ability to reduce their job's priority and, in effect, delay their job's start time by changing this field. This option is usually utilized by users who desire their jobs to execute in a specific order. This field only supports negative numbers with the exception of 0. A 0 user job priority will not delay the job from starting. However, the higher a negative number, the lower a job's priority. For example, a user job priority of -100 will allow more jobs to start before it starts than a user job priority of -10. It should be noted that the user job priority literally lowers the start priority of a job.
Group	Data Dependent	This field will display the group ID under which the job executes.
Generic Attributes	Data Dependent	This field displays a custom attribute attached to the job. Generic attributes are not yet supported in Cluster Manager.
Account	Data Dependent	This field will display the account ID used by the job.
Class/Queue	Data Dependent	This field will display the class/queue ID used by the job.
Quality of Service (QoS)	Data Dependent	This field will display the quality of service (QoS) ID used by the job.

> *Attributes*		
Field	**Displayed**	**Field Information**

> *Attributes*		
Arguments	Data Dependent	Some programs provide users with options. This field allows the user the ability to view those options. A user should consult his/her program documentation to learn about the available options.
By Passed in Queue	Data Dependent	This field displays the number of times another job of a lower priority started before this job.
Input File	Data Dependent	Some scripts/executables/programs/applications require input files to be able to execute. This field allows the user the ability to define those files.
Exclusion Node List	Data Dependent	Often times, users require specific nodes for their applications. This field defines a list of nodes the job cannot execute on.
Executable	Data Dependent	A job consists of a script, executable, program, or application. In order for the job to start, it must know the location of the program. This field allows the user the ability to specify that location. Consult your system administrator for more specific information regarding your program's location.
Flags	Data Dependent	Cluster Manager schedules jobs differently according to their flags. Possible flags are: hold, interactive, restartable, and preemptible.
Initial Working Directory	Data Dependent	Some jobs need to be executed in a specific location on each node. This field allows a user the ability to define that location. By default, the job is executed in the user's home directory. Consult your system administrator for information regarding your home directory.

> *Attributes*		
Master Node	Data Dependent	In a cluster, one specific node is in charge of communication with all the other nodes on the cluster. This node is often referred to as the master node, or the head node. This field will display the name of the master node.
Partition Access List	Data Dependent	This field displays the partitions available for this user. Clusters can be divided into different sections commonly called partitions. Consult your system administrator to learn which partition is the best suited for your job.
Resource Manager Job ID	Data Dependent	All jobs, when created, are given a unique ID by the resource manager. This field displays that ID.
Required Memory	Data Dependent	Some jobs require specific amounts of memory. This field allows a job to request the memory it needs for each node. It should be noted that this field is not the total memory across the entire cluster, but only the memory on each node needed by the job. Workload Manager will start this job only on the nodes that have sufficient memory.
Required Nodes	Data Dependent	A node is a computer consisting of 1 or more processors. A job requires at least 1 processor to execute, and therefore must use at least 1 node.
Required Processors	Data Dependent	All jobs require at least 1 processor. This field displays the number of processors used by this job.
Required Reservations	Data Dependent	A user can specify a reservation for this job. If a reservation is specified, the job will execute only on the nodes that are reserved by the reservation.

> *Attributes*		
Run Priority	Data Dependent	This field is used by jobs that are preemptable to decide which job should be preempted. With some exceptions, the higher a job's priority, the more likely it will be preempted.
System ID	Data Dependent	The system ID is used when multiple resource managers are being used.
Start Count	Data Dependent	This field displays the number of times the job has attempted to start executing.
Start Priority	Data Dependent	This field displays the start priority for the job. With some exceptions, the higher a job's priority over other jobs, the sooner it will begin to execute.
Step ID	Data Dependent	All jobs, when created, are given a unique ID by the resource manager. This field displays that ID.
System Priority	Data Dependent	With some exceptions, the higher a job's priority over other jobs, the sooner it will begin to execute. This field allows an administrator the ability to start jobs immediately. Literally, this field adds 1,000,000,000 plus the administrator priority to the start priority, creating a job with an extremely high priority.
Utilized Memory	Data Dependent	This field displays the amount of memory used by the job during execution.
Utilized Processors	Data Dependent	This field displays the number of processors used by the job during execution.

Time Frame		
Field	**Displayed**	**Field Information**
> Start Time	Data Dependent	This field displays the date and time in which the job started or will start.

Time Frame		
> Duration / Wall Clock	Data Dependent	The *duration* is an estimated time of how long the job will execute. If a user's job requires more time than the specified duration, duration violation policies come into effect. Consult your system administrator for more information regarding these policies. If no duration is specified, a default wall time will be applied. Consult your system administrator for more information regarding your cluster's default wall time.
> Completion Time	Data Dependent	This field displays the time the job finished execution.
> Completed Duration / Used Wall Clock	Data Dependent	This field displays the current execution time of the job.
> Queue Time	Data Dependent	This field displays the amount of time the job waited to start.
> Earliest Start Time	Data Dependent	This field displays the user specified date and time in which the job is available to start. Workload Manager will not start the job until after this specified date and time.
> Required Earliest Start Time	Data Dependent	Some jobs are required to start before a specific time. This field displays the time it has to start before.
> Reservation Start Time	Data Dependent	This field displays the start time for a reservation to which a job is attached. This is only applicable to jobs that were attached to reservations when they were created.

> Statistics		
Field	**Displayed**	**Field Information**

> *Statistics*		
> **Memory Seconds Utilized**	Data Dependent	*Memory seconds utilized* is defined as the total amount of memory used by the job times the number of seconds the memory was used. Users should remember that the value is calculated as a sum total of all the memory on the cluster and not on a per node basis.
> **Dedicated Processor Seconds**	Data Dependent	*Dedicated processor seconds* is defined as the total number of processors reserved by Workload Manager for the job times the number of seconds the processors were reserved. Users should remember that the value is calculated as a sum total of all the processors on the cluster and not on a per node basis.
> **Utilized Processor Seconds**	Data Dependent	Processor seconds utilized is defined as the total number of processors used by the job times the number of seconds the processors were reserved. Users should remember that the value is calculated as a sum total of all the processors on the cluster and not on a per node basis.

Resources		
Field	**Displayed**	**Field Information**
Required Node List	Data Dependent	Often times, users require specific nodes for their applications. If a node list is not specified, the nodes needed for the job are gathered from the nodes field.
Required Allocated Node List	Data Dependent	A node is a computer consisting of 1 or more processors. A job requires at least 1 processor to execute, and therefore must use at least 1 node. The allocated node list is a list of the nodes that the job is using.

Resources		
Required Allocated Partition	Data Dependent	This field displays the required partition for this job. Clusters can be divided into different sections which are commonly called *partitions*. Users can only request one specific partition for their job. Consult your system administrator to learn which partition is the best suited for your job.
Required Node Access	Data Dependent	This field displays the policy that the job uses to select which nodes it can access.
Required Node Set	Data Dependent	This field displays the groupings of nodes this job requires.
Required Preferences	Data Dependent	This field displays the required node preferences for this job.
Required Architecture	Data Dependent	Some jobs require a specific architecture. This field allows a user to view the architecture required by this job.
Required Class	Data Dependent	This field displays the required class/queue for this job.
Required Disk per Task	Data Dependent	A *task* is a group of resources that must all be on the same node. One resource in that group is disk space. This field displays the amount of disk in each task, or group of resources, that the user's job requires.
Required Memory per Task	Data Dependent	A *task* is a group of resources that must all be on the same node. One resource in that group is memory. This field displays the amount of memory in each task, or group of resources, that the user's job requires.
Required Processor per Task	Data Dependent	A *task* is a group of resources that must all be on the same node. One resource in that group is a processor. This field displays the number of processors in each task, or group of resources, that the user's job requires.

Resources		
Required Swap per Task	Data Dependent	A *task* is a group of resources that must all be on the same node. One resource in that group is swap space. This field displays the amount of swap in each task, or group of resources, that the user's job requires.
Required Network	Data Dependent	Some jobs require a specific network. This field allows a user to view the network required by this job.
Required Disk on Node	Data Dependent	Some jobs require specific amounts of disk space. This field allows a user to view the required amount of disk space the job needs on each node. It should be noted that this field is not the total disk across the entire cluster but only the disk space on each node.
Required Features on Node	Data Dependent	Some jobs require a specific feature on a node. A feature is a custom tag attached to a specific list of nodes. This field allows a user to view the required feature for the job. Consult your system administrator for specific information regarding each tag.
Required Memory on Node	Data Dependent	Some jobs require specific amounts of memory. This field allows the user to view the requested amount of memory it needs for each node. It should be noted that this field is not the total memory across the entire cluster but only the memory on each node.
Required Processors on Node	Data Dependent	All jobs require at least 1 processor. This field displays the processors required by this job.

Resources		
Required Swap on Node	Data Dependent	Some jobs require specific amounts of swap space. This field allows a user to view the required swap space the job needs for each node. It should be noted that this field is not the total swap across the entire cluster but only the swap on each node.
Required Operating System on Node	Data Dependent	Some jobs require a specific operating system. This field allows a user to view the operating system required by this job.
Maximum Required Nodes	Data Dependent	A *node* is a computer consisting of 1 or more processors. A job requires at least 1 processor to execute and therefore must use at least 1 node. This field displays the maximum required nodes for the job.
Minimum Required Nodes	Data Dependent	A *node* is a computer consisting of 1 or more processors. A job requires at least 1 processor to execute and therefore must use at least 1 node. This field displays the minimum required nodes for the job.
Maximum Required Tasks	Data Dependent	A *task* is a group of resources that must all be on the same node. This field displays the maximum number of tasks, or groups of resources, the user's job requires on each node.
Minimum Required Tasks	Data Dependent	A *task* is a group of resources that must all be on the same node. This field displays the minimum number of tasks, or groups of resources, the user's job requires on each node.

Resources		
Task Per Node	Data Dependent	A *task* is a group of resources that must all be on the same node. This field displays the number of tasks, or groups of resources, the user's job requires on each node.

Dynamic job attributes will only show jobs that are dynamic and for Moab job templates. The ranged values take a minimum and a maximum value and also show a current value if one exists. A dynamic job will attempt to allocate or deallocate to fit within the specified performance metric ranges.

Dynamic Job Attributes	
Field	**Field Information**
Allocation Delay	*The time in seconds that must take place between any two allocation or deallocaion actions.*
Allocation Size	*The maximum number of nodes that can be allocated or deallocated in any given allocation window.*
Backlog	*The range that the dynamic job's or job template's backlog must be within without reallocation.*
Response Time	*The range that the dynamic job's or job template's response time must be within without reallocation.*
Target Load	*The range that the dynamic job's or job template's load must be within without reallocation.*
Throughput	*The range that the dynamic job's or job template's throughput must be within without reallocation.*
Node Range	*The range that the dynamic job's node count must be within.*
Processor Range	*The range that the dynamic job's processor count must be within.*

2.2.3. List Jobs/Job Templates

Default Categories

Summary:Job ID, Job Name, State, User, Used Wall Clock, Allotted Wall Clock, Nodes, Procs

Credentials:Job ID, State, User, Group, Class, Account, QoS

Time:Job ID, State, Start Time, Used Wall Clock, Completion Time, Submission Time, System Minimal Start Time, Wall Clock, Earliest Start Time, Latest Completion Time

Node Information:Job ID, State, Allocated Node List, Master Node, Node List, Executed Node List

IDs:Job ID, State, Global Job ID, System ID, Step ID, RM Job ID

Required Resources:Job ID, State, Allocated Node List, Partitions, Allocated Nodes, Required Procs, Req. Node Feature, Req. Node Memory

Utilized Resources:Job ID, State, Memory-Seconds Utilized, Processors-Seconds Dedicated, Processor-Seconds Utilized, Utilized Memory, Utilized Processors

Reservation:Job ID, State, Reservation

Resource Manager:Job ID, State, RM Job ID

Executable:Job ID, State, Input File, Executable, Arguments, Initial Working Directory, Executable Type

Priority:Job ID, State, Run Priority, System Priority, User Priority, Start Priority

Diagnostic:Job ID, State,Suspend Duration, Hold, Blocked Reason, Expected State, Bypass

Comments:Job ID, State, Messages

List Jobs Fields

Field	***Category***	***Field Information***
Job ID	All	All jobs, when created, are given a unique ID by Workload Manager. This field displays that unique ID.
Job Name	Summary	A user can attach a custom name to the job to allow him/her to easily identify their job. The name does not change any Workload Manager settings or prioritizations. If a name has been attached it will appear in this field.

List Jobs Fields		
State	All	This field will display the execution status of the job. For example, running, stopped, executing, idle, blocked, etc.
User	Credentials, Summary	This field will display the *user* id under which the job is executing.
Group	Credentials, Summary	This field will display the *group*id under which the job executes.
Class/Queue	Credentials	This field will display the *class/queue*id used by the job.
Account	Credentials	This field will display the *account*id used by the job.
Quality of Service (QoS)	Credentials	This field will display the *quality of service (QoS)* id used by the job.
Start Time	Time, Summary	This field displays the date and time in which the job started.
Used Wall Clock	Time	This field displays the actual execution time of the job. The format is hours:minutes:seconds. The white space indicates the remaining time left before the job reaches it's requested wall clock time. The colored section indicates the amount of wall clock that has been used. A red bar indicates that the job has violated it wall clock limit. Refer to the *Wall Clock* field for the job wall clock.
Completion Time	Time	This field displays the date and time in which the job finished execution.
Submission Time	Time	This field displays the time in which the job was first created/submitted. The format is hours:minutes:seconds. If the exact date is desired moving the mouse over the value will display the exact date of the submission time.

List Jobs Fields		
Suspend Duration	Diagnostics	This field displays the time in which the job was in a suspended state. The format is hours:minutes:seconds.
System Start Time	Time	This field displays the time when the job started.
Execution Eligibility Time	Diagnostics	This field displays the time in which the job was eligible for execution but did not start. The format is hours:minutes:seconds.
Earliest Start Time	>	This field displays the user specified date and time in which the job is available to start. Workload Manager will not start the job until after this specified date and time.
Wall Clock	Time	The duration is an estimated time of how long the job will execute. If a users job requires more time than the specified duration, duration violation policies come into effect. Consult your system administrator for more information regarding these policies. If no duration is specified, a default wall time will be applied. Consult your system administrator for more information regarding your clusters default wall time.
Latest Completion Time	Time	This field displays the date and time in which the job must finish execution.
Allocated Node List	Node Information, Required Resources	A node is a computer consisting of 1 or more processors. A job requires at least 1 processor to execute and therefore must use at least 1 node. The allocated node list is a list of the nodes that the job is using.

List Jobs Fields		
Master Node	Node Information	In a cluster one specific node is in charge of communication with all the other nodes on the cluster. This node is often referred to as the master node or the head node. This field will display the name of the master node.
Node List	Node Information	A node is a computer consisting of 1 or more processors. This field displays the list of nodes that the job requires to execute.
Excluded Node List	Node Information	A node is a computer consisting of 1 or more processors. This field defines the a list of nodes that the job is *not* allowed to use.
Global Job ID	IDs	The *global job id* is used when multiple resource managers are being used.
System ID	IDs	The *system job id* is used when multiple resource managers are being used.
Step ID	IDs	The step id is used by some resource managers to track the job.
Resource Manager Job ID	IDs, Resource Manager	All jobs, when created, are given a unique ID by the resource manager. This field displays that ID.
Input File	Executable	Some scripts/executables/programs/applications required input files to be able to execute. This field allows the user the ability to define those files.

List Jobs Fields		
Executable	Executable	A job consists of a script, executable, program, or application. In order for the job to start it is necessary for it to know the location of the program. This field allows the user the ability to specify that location. Consult your system administrator for more specific information regarding your programs location.
Arguments	Executable	Some programs provide users with options. This field allows the user the ability to view those options. A user should consult his/her program documentation to learn about the available options.
Initial Working Directory	Executable	Some jobs need to be executed in a specific location on each node. This field allows a user the ability to define that location. By default, the job is execute in the user's home directory. Consult your system administrator for information regarding your home directory.
Hold	Diagnostics	A hold can only be placed upon jobs that haven't started. A hold stops or halts a job from running until the user or an administrator releases the hold. If a hold has been placed it will be displayed in this field.
Blocked Reason / Error	Diagnostics	This field will display diagnostic messages relating to the job.
Messages	Comments	This field gives the user the option of adding a comment to the job.
Expected State	Diagnostics	This field displays the execution status of the job that Workload Manager assumes that job is at. For example, running, stopped, executing, idle, blocked, etc.

List Jobs Fields		
By passed in Queue	Diagnostics	This field displays the number of times another job of a lower priority started before this job.
Partitions	Required Resources	Clusters are often divided into different sections. These sections are commonly called partitions. Users can only request one specific partition for their job. Consult your system administrator to learn which partition is the best suited for your job.
Required Quality of Service	>	This field displays the required *quality of service (qos)*for this job.
Required Memory	>	Some jobs require specific amounts of memory. This field allows a job to request the memory it needs for each node. It should be noted that this field is not the total memory across the entire cluster but only the memory on each node needed by the job. Workload Manager will start this job only on the nodes that have sufficient memory.
Nodes	Required Resources, Summary	A node is a computer consisting of 1 or more processors. A job requires at least 1 processor to execute and therefore must use at least 1 node. This fields displays the number of nodes used by the job.
Processors	Required Resources, Summary	All jobs requires at least 1 processor. This field displays the number of processors used by this job.
Required Nodes	>	A node is a computer consisting of 1 or more processors. A job requires at least 1 processor to execute and therefore must use at least 1 node.

List Jobs Fields		
Reservation	Reservation	A user can specify a reservation for this job. If a reservation is specified the job will execute only on the nodes that are reserved by the reservation.
Run Priority	Priority	This field is used by jobs that are preemptible to decide which job should be preempted. With a few exceptions, the higher a job's priority, the more likely it will be preempted.
Start Count	>	This field displays the number of times the job has attempted to start executing.
System Priority	Priority	With a few exceptions, a job with a high priority will begin sooner depending on how much greater its priority is to other jobs. This field allows an administrator the ability to start jobs immediately. Literally, this field adds 1,000,000,000 plus the administrator priority to the start priority creating a job with an extremely high priority.

List Jobs Fields		
User Job Priority	Priority	With a few exceptions, a job with a high priority will begin sooner depending on how much greater its priority is to other jobs. A user has the ability to reduce their job's priority and, in effect, delay their job's start time by changing this field. This option is usually utilized by users who desire their jobs to execute in a specific order. This field only supports negative numbers with the exception of 0. A 0 user job priority will not delay the job from starting. However, a job's priority will decrease as the priority number decreases. For example, a user job delay priority of -100 will allow more jobs to start before it starts than a user job priority of -10. It should be noted that the user job priority literally lowers the start priority of a job.
Start Priority	Priority	This field displays the start priority for the job. With a few exceptions, a job with a high priority will begin sooner depending on how much greater its priority is to other jobs.
Memory-Seconds Utilized	Utilized Resources	Memory seconds utilized is defined as the total amount of memory *used* by the job times the number of seconds the memory was used. Users should remember that the value is calculated as a sum total of all the memory on the cluster and not on a per node basis.

List Jobs Fields		
Processor-Seconds Dedicated	Utilized Resources	Processor seconds dedicated is defined as the total number of processors *reserved* by Workload Manager for the job times the number of seconds the processors were reserved. Users should remember that the value is calculated as a sum total of all the processors on the cluster and not on a per node basis.
Processor-Seconds Utilized	Utilized Resources	Processor seconds utilized is defined as the total number of processors *used* by the job times the number of seconds the processors were reserved. Users should remember that the value is calculated as a sum total of all the processors on the cluster and not on a per node basis.
Utilized Memory	Utilized Resources	This field displays the amount of memory used by the job during execution.
Utilized Processors	Utilized Resources	This field displays the number of processors used by the job during execution.
Flags	>	Cluster Manager schedules jobs differently according to their flags. Possible flags are hold, interactive, restartable, and preemptible. Refer to the create job documentation for definitions of the flags.
Generic Attribute	>	This field displays a custom attribute attached to the job. Generic attributes are not supported in Cluster Manager yet.
Required Allocated Node List	>	A node is a computer consisting of 1 or more processors. A job requires at least 1 processor to execute and therefore must use at least 1 node. The allocated node list is a list of the nodes that the job is using.

List Jobs Fields		
Required Allocated Partition	>	Clusters are often divided into different sections. These sections are commonly called partitions. Users can only request one specific partition for their job. Consult your system administrator to learn which partition is the best suited for your job.
Required Node Access	>	This field displays the policy that job uses to select which nodes it can access
Required Node Set	>	This field displays the groupings of nodes this job requires.
Required Preferences	>	This field displays the required node preferences for this job.
Required Architecture	>	Some jobs require a specific node architecture. This field allows a user to view the architecture required by this job.
Required Class/Queue	>	This field displays the required *class/queue*for this job.
Required Disk Per Task	>	A task is a group of resources that must all be on the same node. One resource in that group is disk space. This field displays the amount of disk in each task, or group of resources, that the users job requires.
Required Memory Per Task	>	A task is a group of resources that must all be on the same node. One resource in that group is memory. This field displays the amount of memory in each task, or group of resources, that the users job requires.
Required Network	>	Some jobs require a specific network. This field allows a user to view the network required by this job.

List Jobs Fields		
Required Disk On Node	>	Some jobs require specific amounts of disk space. This field allows a user to view the required amount of disk space the job needs on each node. It should be noted that this field is not the total disk across the entire cluster but only the disk space on each node.
Required Node Features	Required Resources	Some jobs require a specific feature on a node. A feature is a custom tag attached to a specific list of nodes. This field allows a user to view the required feature for the job. Consult your system administrator for specific information regarding each tag.
Required Node Memory	Required Resources	Some jobs require specific amounts of memory. This field allows the user to view the requested amount of memory it needs for each node. It should be noted that this field is not the total memory across the entire cluster but only the memory on each node.
Required Processors On Node	>	All jobs require at least 1 processor. This field displays the processors required by this job.
Required Swap On Node	>	Some jobs require specific amounts of swap space. This field allows a user to view the required swap space the job needs for each node. It should be noted that this field is not the total swap across the entire cluster but only the swap on each node.
Required Operating System	>	Some jobs require a specific operating system. This field allows a user to view the operating system required by this job.

List Jobs Fields		
Required Partition	>	Clusters are often divided into different sections. These sections are commonly called partitions. Users can only request one specific partition for their job. Consult your system administrator to learn which partition is the best suited for your job.
Required Processors Per Task	>	A task is a group of resources that must all be on the same node. One resource in that group is a processor. This field displays the number of processors in each task, or group of resources, that the users job requires.
Required Swap Per Task	>	A task is a group of resources that must all be on the same node. One resource in that group is swap space. This field displays the amount of swap in each task, or group of resources, that the users job requires.
Required Maximum Nodes	>	A node is a computer consisting of 1 or more processors. This field displays the maximum number of nodes required for the job to execute.
Required Minimum Nodes	>	A node is a computer consisting of 1 or more processors. This field displays the minimum number of nodes required for the job to execute.
Required Maximum Task Count	>	A task is a group of resources that must all be on the same node. This field displays the maximum number of tasks, or groups of resources, the users job requires on each node.
Required Minimum Task Count	>	A task is a group of resources that must all be on the same node. This field displays the minimum number of tasks, or groups of resources, the users job requires on each node.

List Jobs Fields		
Required Tasks Per Nodes	>	A task is a group of resources that must all be on the same node. This field displays the number of tasks, or groups of resources, the users job requires on each node.
Executable Type	>	If type is known, this will display if an executable is a binary executable or script executable.
Bypass	>	This displays the number of times the job has been 'bypassed' by a lower priority job via backfill.

2.2.4. Job Timeline

The Job Timeline window displays each executing job. On the left side of the Now line is the amount of a job that has completed while the right side is the remaining execution time.

2.2.5. Job Outlines

Job outlines are settings saved from the create/submit job window. This window allows job outlines to be saved, deleted, or opened either locally on the machine that Moab Cluster Manager is running on, or remotely on the machine that Moab Workload Manager is running on.

Outline Information - This section displays information of the currently loaded job outline.	
Field	**Field Information**
Outline Name	This field allows a user to input or change the saved file name of the outline.
Name	This field displays the custom name for the job outline.
Script	This field displays the script that is used in the outline.

2.2.5.1. Outline Windows

Local Job Outlines - Outlines found on the machine that Moab Cluster Manager is running on.

Personal Remote Outlines - Outlines found in a personal directory on the machine running Moab

Workload Manager.

Shared Remote Outlines - Outlines found in a shared remote directory on the machine running Moab Workload Manager.

Outline Directories - Directories where outlines can be found.

Outline Operations	
Field	**Field Information**
Load	This button will get the selected outline and insert the information in the outline information fields
Delete	This button will delete the selected outline.

2.2.6. Dynamic Job Allocation

Dynamic job allocation allows a user with mjobctl privileges to manually allocate or deallocate nodes for a dynamic job. This allows complete control over how many nodes a dynamic job has. Keep in mind that if performance metric ranges are specified (such as TARGETLOAD, TARGETBACKLOG, etc.) that the dynamic job may reallocate or deallocate nodes that were just modified to meet their metrics.

If this window was accessed from a node based window (such as "Visual Cluster" or "List/Modify Nodes"), the node list will be locked down. The jobs will be filtered such that they make sense in the add or release action. Similarly, if this window was accessed from a job based window (such as "List/Modify Jobs"), the job list will be locked down. The nodes will be filtered such that they make sense in the add or release action.

2.3. Reservations

2.3.1. Create Reservation

Summary:

A *reservation* sets apart resources during a particular time frame for a particular owner. Reservations usually reserve resources such as nodes or processors on the cluster. The create reservation window allows you to define what resources a reservation requires, as well as the time frame for the reservation.

In addition to being able to create a basic reservation, this window also allows you to create a recurring reservation. A *recurring reservation* , also referred to as a standing reservation or a reservation

generator/creator, provides the user with the option of having reservations automatically created according to a desired time frame. For example, If a user wanted a reservation to be created every Tuesday and Thursday starting at 11 am and ending at 4 pm, a recurring reservation would fulfill this need. It should be noted that a recurring reservation could potentially not be able to create a reservation if the resources are already dedicated to another reservation or job. To reduce the possibilities of this occurring the day/week depth field can be increased so Workload Manager will create the reservations further in advance.

Panels Containing Required Parameters:

Reservation Information
This panel allows the user to specify the reservation name and owner.

Field	*Required*	*Additional Information*
Reservation Name	Required	This field allows a user to create a name for the reservation. Workload Manager appends a numerical value to the end of the Reservation Name which allows users the ability to enter duplicate Reservation Name's without affecting a previous reservation.
Owner	Optional	An *owner* is a user, group, account, class, or quality of service. A reservation can reserve only the resources that the owner has access to. This field allows a user to select the owner of the reservation.

Access Control List
This panel allows the user to specify what credentials have permission to access the reservation. At least one credential is required to be in the Access Control List for a reservation, otherwise it would not be very useful. The user may select from 5 different type of credentials - users, groups, accounts, classes, and quality-of-services. Any of the credentials in the Access Control List have permission to use the reservation.

Button	*Additional Information*	
Add	Pops up a wndow that allows the user to select credentials of a specific type to add to the Access Control List	
Clear	Removes all credentials fromthe Access Control List	

Resources **Allows the user to specify what resources will be**		
Field	***Required***	***Additional Information***
Host List / Host Expression	Required if Task Count isn't defined.	Often times, users require specific nodes for their applications. A list of nodes required by the user is commonly called a *host list*. If a host list is not specified, the task count must be specified.
Search Resources	Optional	This button displays a table allowing the user the ability to search for available resources. Consult the search resources documentation for more specific information.
Task Count	Required if Host List isn't populated	A *task* is a group of resources that must all be on the same node. This field defines how many groups of resources will be required to create this reservation. If the task count is not specified the host list must be specified.

Resources **Allows the user to specify what resources will be**		
Tasks	Optional	A *task* is a group of resources that must all be on the same node. This button displays a window allowing a user to define what those resources will be in a task. • Memory Per Task - This field provides the user with the option of requesting the amount of memory, in megabytes, a job needs on a node. • Swap Per Task - This field provides the user with the option of requesting the amount of swap, in megabytes, a job needs on a node. • Disk Per Task - This field provides the user with the option of requesting the amount of disk space, in megabytes, a job needs on a node. • Processors Per Task - This field provides the user with the option of requesting the number of processors a job needs on a node.

Reservation Time Frame Tabs

Once - Basic Reservation Creation **One basic reservation will be created for the desired start and end time.**		
Field	***Required***	***Additional Information***
Start Time	Required	This field allows the user to select the day and time when the reservation begins.

Once - Basic Reservation Creation **One basic reservation will be created for the desired start and end time.**		
End Time	Required	This field allows the user to select the day and time when the reservation ends.

Daily Recurring Reservation **A recurring reservation will be initialized to automatically create a reservation on the desired days at the desired start and end time. The recurring reservation will continually generate/create new reservations. The *day depth* is used to decide when a reservation is created. For example, if a reservation starts 4 days from now and the *day depth* is set to 2 days the reservation will not be created for 2 more days.**		
Field	***Required***	***Additional Information***
Days	Required	This field allows the user to select the days of the week when a reservation is created.
Start Time	Required	This field allows the user to select the time of day when the reservation begins.
End Time	Required	This field allows the user to select the time of day when the reservation ends.
Day Depth	Optional	This field allows the user to specify how many days in advance Workload Manager should create recurring reservations.

Weekly Recurring Reservation **A recurring reservation will be initialized to automatically create a reservation for the week starting from the desired start day and ending on the desired end day of the week. The recurring reservation will continually generate/create new reservations. The *week depth* is used to decide when a reservation is created. For example, if a reservation starts 4 weeks from now and the *week depth* is set to 2 weeks the reservation will not be created for 2 more weeks.**		
Field	***Required***	***Additional Information***
Start Day	Required	This field allows the user to select the day and time in the week when the reservation begins.

Weekly Recurring Reservation **A recurring reservation will be initialized to automatically create a reservation for the week starting from the desired start day and ending on the desired end day of the week. The recurring reservation will continually generate/create new reservations. The *week depth* is used to decide when a reservation is created. For example, if a reservation starts 4 weeks from now and the *week depth* is set to 2 weeks the reservation will not be created for 2 more weeks.**		
End Day	Required	This field allows the user to select the day and time in the week when the reservation ends.
Week Depth	Optional	This field allows the user to specify how many weeks in advance Workload Manager should create recurring reservations.

Infinite Reservation A reservation will be created that will continue indefinitely. No start or end time is required.

Panels Containing Advanced Options:

Option Tabs

Single Reservation Options		
Field	***Required***	***Additional Information***
Exclusive	Optional	The exclusive option allows only this reservation, and no other reservation, access to the requested resources.

Recurring Reservation Options		
Field	***Required***	***Additional Information***
Single Use	Optional	The single use option allows only one job to run in this reservation. Once that job has finished execution, the basic reservation, not the recurring reservation, will be canceled.
By Name	Optional	Only jobs that request this reservation will be allowed to execute within it.

Recurring Reservation Options		
Owner Preempt	Optional	This option allows jobs that are running inside of this reservation to be preempted by jobs owned by the same owner as this reservation.
Space Flex	Optional	The *space flex* option gives Workload Manager permission to alter the number of requested resources for this reservation.
Time Flex	Optional	The *time flex* option gives Workload Manager permission to alter the time frame for this reservation. It should be noted that the space flex option must be enabled if time flex is desired.

Event Triggers		
Button	***Required***	***Additional Information***
CreateTrigger 1-6	Optional	This field allows the user the ability to attach triggers to a reservation.

Grid Sandboxing		
Field	***Required***	***Additional Information***
Allow grid sandboxing	Optional	By default, this allows only the resources in the recurring reservation to be visible to grid peers.
Cluster List	Optional	List of clusters who have access to the grid sandbox.

Misc. Options		
Field	***Required***	***Additional Information***
Partition	Optional	Clusters can be divided into different sections. These sections are commonly called *partitions*. Users can only request one specific partition for their reservation. Consult with your system administrator to learn which partition is best suited for your reservation.

Misc. Options		
Node Features	Optional	Some jobs require a specific feature on a node. A *node feature* is a custom tag attached to a specific list of nodes. Consult your system administrator for specific information regarding each tag.

2.3.2. Modify Reservation

Summary

A *reservation* is a time frame on the cluster reserved for a particular need. Reservations usually reserve resources such as nodes or processors on the cluster. The modify reservation window allows you to view and modify existing reservations.

Basic Information		
Field	**Displayed**	**Field Information**
Reservation Name	Always	This field allows a user to create a name for the reservation. Workload Manager appends a numerical value to the end of the Reservation Name allowing users the ability to enter duplicate Reservation Name's without affecting any other reservation.
Reservation Owner	Always	An *owner* is a user, group, account, class, or quality of service. A reservation can reserve only the resources that the owner has access to. This field displays the owner of the reservation. If the reservation is a job reservation this field will be blank.
Global ID	Data Dependent	This field only displays information when multiple resource managers are present.
Messages	Data Dependent	This field allows a user the option of adding a message or comment to a reservation.

Basic Information		
Type	Data Dependent	This field displays whether a reservation is a user or job reservation.
Sub Type	Data Dependent	This is the type of reservation. Some examples of the available types are grid, standing reservation, user, maintenance, etc.
Trigger	Data Dependent	This field displays information about any trigger that is attached to the reservation.
Statistics	Data Dependent	This field displays the percentage of processors seconds reserved by the reservation that were used by a job or multiple jobs.

Credentials		
Field	**Displayed**	**Field Information**
User	Always	This field will display the user ID used by the reservation. If this field is empty no user can directly access this reservation.
Group	Always	This field will display the group ID used by the reservation. If this field is empty no group can directly access this reservation.
Account	Always	This field will display the account ID used by the reservation. If this field is empty no account can directly access this reservation.
Class	Always	This field will display the class/queue ID used by the reservation. If this field is empty no class can directly access this reservation.
Quality of Service (QoS)	Always	This field will display the quality of service (QoS) ID used by the reservation. If this field is empty no quality of service can directly access this reservation.

Cluster Information

Cluster Information		
Field	**Displayed**	**Field Information**
Flags	Data Dependent	Cluster Manager schedules reservations differently according to their flags. This field displays the reservation flags.
Node List Regular Expression	Data Dependent	This field displays a list of nodes required by the job to execute. The list of nodes is a regular expression. A *node* is a computer consisting of 1 or more processors.
Job ID	Data Dependent	This field displays the job ID of a job reservation.
Processor	Data Dependent	This field displays the number of processors used by a job for a job reservation.
Processor-Seconds	Data Dependent	This field displays the number of processor-seconds used by a job for a job reservation.
Max Tasks	Data Dependent	This fields displays the maximum number of tasks a reservation can use. A *task* is a group of resources that must all be on the same node.
Required Feature List	Data Dependent	A *feature* is a custom attribute attached to a node. This field displays the features required to be on a node for the reservation to reserve the node.
Required Feature Policy	Data Dependent	This field displays the policy that the reservation will use to select the features.
Required Node Count	Data Dependent	This field displays the number of nodes required by the reservation. A *node* is a computer consisting of 1 or more processors.
Required Node List	Data Dependent	This field displays a list of nodes required by the reservation. A *node* is a computer consisting of 1 or more processors.
Required Task Count	Data Dependent	This field displays the number of processors required by the reservation.

Cluster Information		
Resources	Data Dependent	This field displays what type of resource is reserved by the reservation.
Spec Name	Data Dependent	This field is displays information for multiple resource managers.
Task Count	Data Dependent	This field displays how many groups of resources will be required to create this reservation. A *task* is a group of resources that must all be on the same node.

Time Frame		
Field	**Displayed**	**Field Information**
Start Time	Data Dependent	Some jobs require a specific amount of time before they should be allowed to start. This field allows the user to define the earliest time the job can start. By default, a job may start as soon as resources become available.
Duration	Data Dependent	The *duration* is an estimated time of how long the job will take to execute. If a user's job requires more time than the specified duration, duration violation policies come into effect. Consult your system administrator for more information regarding these policies. If no duration is specified, a default wall time will be applied. Consult your system administrator for more information regarding your cluster's default wall time.
End Time	Data Dependent	This field displays the time at which the reservation will finish.

Node List		
Field	**Displayed**	**Field Information**

Node List		
Nodes	Data Dependent	Each button displayed represents a node that the reservation has reserved. When the button is selected the view/modify node window will appear containing information about the node.

2.3.3. List Reservations

Summary

A *reservation* is a time frame on the cluster reserved for a particular needed. Reservations usually reserve resources such as nodes or processors on the cluster. Reservations are created either by a user or by a job. A reservation created by a user is called a *user reservation* while areservationcreated by a job is called a *job reservation*. All executing jobs have reservations.

List Reservation Fields		
Field	*Category*	*Field Information*
Name	All	This field allows a user to create a name for the reservation. Workload Manager appends a numerical value to the end of the Reservation ID allowing users the ability to enter duplicate Reservation IDs without affecting any other reservation.
Type	All	This field displays whether a reservation is a user or job reservation.
User	Summary, Credentials	This field displays which users will be able to access this reservation. If this field is blank no users have been given access to the reservation.
Group	Summary, Credentials	This field displays which groups will be able to access this reservation. If this field is blank no groups have been given access to the reservation.

List Reservation Fields		
Account	Summary, Credentials	This field displays which accounts will be able to access this reservation. If this field is blank no accounts have been given access to the reservation.
Class	Summary, Credentials	This field displays which classes/queues will be able to access this reservation. If this field is blank no classes/queues have been given access to the reservation.
Quality of Service (QoS)	Summary, Credentials	This field displays which quality of service (QoS) will be able to access this reservation. If this field is blank no qualities of service (QoS) have been given access to the reservation.
Start Time	Summary, Time	If the reservation will start in less than 12 hours, the value display is in the format of hours:minutes:seconds, where a negative value indicates that the reservation will start in that many hours:minutes:seconds. A positive value indicates that the reservation started that many hours:minutes:seconds ago. Resting the mouse over the value will display the exact date that the reservation started or will start. The colored bar shows the percentage of the reservation that has completed. The white space indicates the remaining reservation time.

List Reservation Fields		
End Time	Summary, Time	If the reservation will end in less than 12 hours, the value display is in the format of hours:minutes:seconds, where a negative value indicates that the reservation ended that many hours:minutes:seconds ago. A positive value indicates that the reservation will end in that many hours:minutes:seconds. Resting the mouse over the value will display the exact date that the reservation ended or will end. An end time that is years in the future often indicates that the reservation was created without any end time specified and Workload Manager inserted a default end time.
Allocated Node List	Nodes	A *node* is a computer consisting of 1 or more processors. The allocated node list is a list of the nodes that the reservation is using.
Duration	Time	The duration is an estimated time of how long the job will execute. The format used is days:hours:minutes:seconds.
Flags		Cluster Manager schedules reservations differently according to their flags. This field displays the reservation flags.
Global ID	Identification	This field only displays information when multiple resource managers are present.
Node List Regular Expression	Nodes	A *node* is a computer consisting of 1 or more processors. This field displays a list of nodes required by the job to execute. The list of nodes is a regular expression.
Maximum Tasks	Tasks	This fields displays the maximum number of processors a reservation can use.

List Reservation Fields		
Messages		This field allows a user the option of adding a message or comment to a reservation.
Node Count	Nodes	A *node* is a computer consisting of 1 or more processors. This field displays the number of nodes used by the reservation.
Node List	Nodes	A *node* is a computer consisting of 1 or more processors. This field displays a list of nodes being used by the reservation.
Node Set Policy	Nodes	This field displays the policy that the reservation will use to select the nodes.
Owner	Identification	This field displays the owner of the reservation. A reservation can reserve only the resources that the owner has access to. An *owner* is a user, group, account, class, or quality of service.
Partition		Clusters can be divided into different sections. These sections are commonly called *partitions*. Users can only request one specific partition for their reservation. Consult your system administrator to learn which partition is the best suited for your reservation.
Required Feature List	Required Resources	A feature is a custom attribute attached to a node. This field displays the features required to be on a node for the reservation to reserve the node.
Required Feature Policy	Required Resources	This field displays the policy that the reservation will use to select the features.
Required Node Count	Required Resources	A *node* is a computer consisting of 1 or more processors. This field displays the number of nodes required by the reservation.

List Reservation Fields		
Required Node List	Required Resources	A *node* is a computer consisting of 1 or more processors. This field displays a list of nodes required by the reservation.
Required Task Count	Required Resources, Tasks	This field displays the number of processors required by the reservation.
Resources		This field displays what type of resource is reserved by the reservation.
Statistics	Statistics	This field displays statistical information relating to the reservation.
Specification Name		This field displays information for multiple resource managers.
Sub Type		This displays the type of reservation. Some examples of the available types are grid, standing reservation, user, maintenance, etc.
Task Count	Tasks	A *task* is a group of resources that must all be on the same node. This field displays how many groups of resources will be required to create this reservation.
Trigger		This field displays information about any trigger that is attached to the reservation.
Processors		This field displays the number of processors used by the reservation.
Time line		This displays the reservation time lines. The green bar indicates the used amount of the reservation while the blue bar indicates the remaining amount of the reservation. The display options on the left side allow a user change how much of the time line is displayed.

Default Category Settings

Summary:Name, Type, User, Group, Account, Class, Quality of Service (QoS), Start Time, End Time, Duration

Credentials: Name, Type, User, Group, Account, Class, Quality of Service (QoS)

Time: Name, Type, Start Time, End Time, Duration

Resources: Name, Type, Partition, Resources

Required Resources: Name, Type, Required Feature List, Required Node count, Required Node List

Flags: Name, Type, Flags

Nodes: Name, Type, Allocated Node List, Node Expression, Node Count, Node List, Node Set Policy

Statistics: Name, Type, Statistics

Comments: Name, Type, Messages

Tasks: Name, Type, Maximum Tasks, Required Task Count, Task Count

Identification: Name, Type, Global ID, Owner

Trigger: Name, Type, Trigger

2.3.4. List Recurring Reservations

Summary

A *recurring reservation*, also referred to as a *standing reservation* or a *reservation generator*, creates reservations according to user-defined settings. To choose which fields you would like to view in the chart, click on the customize table columns icon, which is the second icon to the left. To add fields, create a check mark by clicking in the box to the left of the field you would like to view. To remove fields, click on the checked box.

List Reoccurring Reservations	
Field	*Field Information*
ID	This field displays the reservation generator's ID.

List Reoccurring Reservations	
Host List	The *host list* is a list of the nodes that the reservation is using. A *node* is a computer consisting of 1 or more processors.
Owner	A reservation generator can reserve only the resources to which the owner has access. This field displays the owner of the reservation generator. An *owner* is a user, group, account, class, or quality of service.
User	This field displays which users will be able to access the created reservation.
Group	This field displays which groups will be able to access the created reservation.
Account	This field displays which accounts will be able to access the created reservation.
Class	This field displays which classes /queues will be able to access the created reservation.
Quality of Service (QoS)	This field displays which quality of service will be able to access the created reservation.
Period	This field displays the interval in which the reservations will be created. The display options are daily, weekly or infinitely.
Days	This field displays the days in the week when the reservations will start. This field will only display information if the period is set to daily.
Disable	This field displays whether a particular reservation generator setting has been disabled.

2.3.5. Reservation Timeline

Summary

The Reservation Timeline window displays each Reservation. On the left side of the Now line is the amount of the Reservation that has been used, while the right side is the remaining amount of the reservation.

2.3.6. Reservation Calendar

Summary:

The reservation calendar displays reservations color coded by AAccount. If a reservation does not have

an AAccount, it is shown in gray. The height of a reservation indicates the number of processors it needs. Thus, tall reservations require more processors. Detailed information about the reservation can be seen upon mouseover. To modify a reservation, simply click on it and a reservation modification window will appear.

2.4. Triggers

2.4.1. List Triggers

Summary:

Workload Manager can launch events or triggers based on certain events. For example, an administrator may want an email sent when the reservation usage falls below a certain percentage, or a user may want to launch an evaluation script 5 minutes before his or her job is scheduled for completion.

List Triggers		
Field	**Category**	**Field Summary**
Trigger ID		This field displays the unique ID assigned to the trigger by Workload Manager.
Trigger State		This field displays the execution status of the trigger. If the state is Idle, the trigger is waiting to execute. If the state is Active, the trigger is executing. Once the trigger has executed, the state displayed will be Successful or Failure, depending on the outcome of the trigger action.
Resource ID		This field displays the ID of the job, reservation or node to which the trigger is attached.
Resource Type		This field displays whether the trigger is attached to a job, reservation, node or the scheduler.

List Triggers		
Resource Event		This field displays the event that must occur for the trigger to execute. The possible events are: when the resource is created, when the resource starts, when the resource ends, or when a failure occurs in the resource.
Trigger Action		This field displays the type of trigger action that will occur when the trigger is executed. The possible trigger types are: cancel the resource the trigger is attached to, email the administrator or execute a script/application/program/executable.
Script		This field displays the script/application/program/executable that will be executed when the trigger is executed.
Seconds Offset		This field displays the number of seconds after a resource event occurs the trigger will execute. If this value is negative, the trigger will execute that many seconds before the resource event occurs.
Flags		This field displays which flags have been set for this trigger.
Threshold		This field displays the reservation usage threshold for this trigger. If the reservation falls below the displayed usage, the trigger will execute.
Estimated Start Time		This field displays the date and time when the trigger will execute if it is possible to calculate it. For example, reservation end times and job completion times can be calculated. Node or scheduler failures can not be calculated ahead of time.

List Triggers		
Actual Start Time		This field displays the date and time when the trigger started. This field is only populated after the trigger has been executed.
Messages		This field displays the status information indicating possible failures or unexpected conditions.
Output File		This field displays the location of the file containing all the trigger output messages.
Error File		This field displays the location of the file containing all the trigger error messages.

Chapter 3. Resources

3.1. Resources Overview

The resources category gives adminstrators the ability to view, modify, and set policies and attributes while effectively diagnosing various system resources.

3.2. Moab Workload Manager

3.2.1. Control Panel

This window provides a control center for the basic operations of Workload Manager.

Control Panel

- Pause/Resume- The "Pause" button will stop Workload Manager from scheduling any new jobs but will not turn Workload Manager off. The "Resume" button, which replaces the pause button when Workload Manager is paused, will allow Workload Manager to begin scheduling jobs again.
- Shutdown- The "Shutdown" button will turn Workload Manager off. Please note that Workload Manager cannot be restarted from Cluster Manager.
- Stop Iteration- The "Stop Iteration" button will cause Workload Manager to stop scheduling once it reaches the iteration defined in the field. The iteration is defined as the cycle that Workload Manager is currently on. When Workload Manager starts its cycle is 0. Approximately every 30 seconds Workload Manager increases the cycle by one.

Operating Modes

- **Normal Mode** - This mode is the fully operational mode for Workload Manager.
- **Simulation Mode** - This mode is used for observing a virtual cluster, as well as virtual jobs to observe how Workload Manager schedules.
- **Interactive Mode -** This mode halts and waits for user input before continuing to operate.
- **Monitor Mode** - This mode is used to connect Workload Manager to a live resource manager and monitor the entire cluster. However, in this mode Workload Manager can not change any resource manager decisions. Workload Manager only observes the system.

3.2.2. Log Settings

Summary

Logging is defined as recording error, diagnostic, and informational messages to a file. This window allows an administrator to configure the logging that occurs in Workload Manager.

Options

Field	Additional Information
Log Level	This field allows an administrator the option of specifying the amount of data recorded in the log files. A value of 1 means almost no data is recorded while a value of 9 means all the data is recorded. Each value increment means that approximately double the amount of data is logged to the log files. The default log level is 3.

Field	Additional Information
Log Facilities	This field determines what is recorded in the log file. • fCore - This option records Workload Manager core messages. • fSched - This option records messages that involve the scheduler. • fSock - This option records messages that involve the socket communication. • fUI - This option records messages that involve the user interface. • fLL - This option records messages that come from LoadLeveler Resource Manager. • fRM - This option records resource manager messages. • fSDR - This option records messages that involve system data repository. • fCONFIG - This option records messages that involve the configuration file. • fSTAT - This option records messages that involve statistics. • fSIM - This option records messages that occur during the simulation operating mode. • fSTRUCT - This option records messages that involve Workload Manager's structure. • fFS - This option records messages that involve fairshare. • fCKPT - This option records messages that involve the checkpoint file. • fBANK - This option records messages that involve QBank. • fPBS - This option records messages that come from the Torque/OpenPBS Resource Manager. • fWIKI - This option records messages that involve WIKI. • fALL - This option records all the events that occur.

Field	Additional Information
Log Directory	This field allows an administrator the option of specifying the directory in which log files will be maintained.
Log File	This field allows an administrator the option of specifying the name of the Workload Manager log file.
Log File Max Size	This field allows an administrator the option of specifying the maximum allowed size (in bytes) of the log file before it will be *rolled*.
Log File Roll Depth	When a log file reaches it's maximum size it is *rolled* or renamed to another filename and a new log file is created using the original file name. This field allows an administrator the option of defining the number of renamed files Workload Manager should maintain.

3.2.3. System Settings

Summary

This window provides an administrator the option of changing numerous Workload Manager settings.

Options

Field	Additional Information
Name	This field allows an administrator to name the cluster. The name is only available for administrator convenience and is not used by Workload Manager.
Host	This field allows an administrator to define the host name that the Workload Manager subcomponents or clients will use to connect to Workload Manager.
Port	This field allows an administrator to define the port that the Workload Manager subcomponents or clients will use to connect to Workload Manager.
Home Directory	This field allows an administrator to define the directory where Workload Manager's configuration, statistics, and log files are located.

Field	Additional Information
Feedback Program	This field allows an administrator to define a program that will be run at the completion of each job. Usually the program is used to contact the user through email informing him/her that the job completed execution.
Notify Program	This field allows an administrator to define a program that will be run when messages or alerts occur in Workload Manager.
Node Purge Time Limit	This field allows an administrator to define the amount of time Workload Manager will keep track of a node which is no longer reported by the resource manager. This value should be increased when using a resource manager that often looses information about a node due to internal failures.
Resource Manager Poll Interval	This field is the time in between which Workload Manager will communicate with the resource manager.
Job Purge Time Limit	This field allows an administrator to define the amount of time Workload Manager will keep track of a job which is no longer reported by the resource manager. This value should be increased when using a resource manager that often loses information about a job due to internal failures.
Charge Metric	This field allows an administrator to specify how quality of service charging should occur. 1. DEBITALLCPU - This policy will charge according to the number of processors used. 2. DEBITALLPE - This policy will charge according to the number of processors used times the number of processor equivalent* nodes used. 3. DEBITSUCCESSFULWC - This policy will charge jobs that successfully completed according to the number of hours they were on the cluster. 4. DEBITSUCCESSFULCPU - This policy will charge jobs that successfully completed according to the number of processors used. 5. DEBITSUCCESSFULPE - This policy will charge jobs that successfully completed according to the number of processor equivalent* nodes they used.

Field	Additional Information
Charge Rate Policy	This field allows an administrator to specify how the quality of service charging should occur. 1. QOSREQ - This policy will charge based upon the quality of service requested. 2. QOSDEL - This policy will charge based upon the quality of service dedicated or given.
Service Provisioning	This field allows interaction with provisioning of resources via a provisioning manager. Provisioning is the process of modifying resources to meet existing needs.

*Processor equivalence is a relative measure of how much of a node is taken by a job, even if only one type of node resource is requested. For example, if a job requires 1 processor and 1 GB of memory, and it is running on a 4 processor node with 1GB of memory, the PE of the job is 4. All of the processors are considered to be taken because the first job is using all of the memory, which prevents any other job from running on that node.

3.2.4. Simulation Settings

Summary

Simulation settings are only applicable if Workload Manager is operating in simulation mode. Simulation is used to virtually observe a cluster and how Workload Manager will schedule jobs across the cluster.

Field		Additional Information
Workload Trace		A *Workload Trac* information relat that Workload M simulate schedul requires the locat Workload Trace f
Resource Trace		A *Resource Trac* contains the infor related to the nod Workload Manag simulate schedul This field require location of the R Trace file.

Field		Additional Information
Simulation Job Policy		This field allows administrator the of specifyingwhe Workload Manag add new jobs fro *Workload Trace* f be scheduled.
Initial Queue depth		This field allows administrator to s how many jobs th Workload Manag add to the job qu the Workload Tra
Time Ratio		This field allows administrator to s how fast Worklo Manager will sim job execution. If value is set to 1, will execute at no speed. If this valu to 2, the jobs wil execute at double normal speed. Et

Field		Additional Information
Auto Shutdown		This field allows administrator to s whether Workloa Manager will shu once all the jobs the Workload Tra have been simula

3.2.5. Statistics Settings

Summary

Credential statistics are disabled by default although cluster-wide statistics are always enabled.

Enable Credentials Statistics

This section allows an administrator to enable or disable user, group, account, class, or quality of service statistics. Because statistics increase Workload Manager's memory usage, an administrator can decrease the memory footprint of Workload Manager by disabling credential statistics.

Number of intervals in each day

Workload Manager combines statistics into intervals. The Daily Statistical Count allows an administrator the option of increasing or decreasing the number of intervals in each day. A higher amount of intervals creates more precise statistics, but Workload Manager uses more memory when intervals are higher

3.2.6. High Availability

Summary

High availability provides a backup Workload Manager in the unlikely situation of a failure.

High availability allows an administrator to state the connection information for the primary Workload Manager and a backup secondary Workload Manager.

The fields for both the primary and secondary Workload Manager are as follows:

- Status - This field displays information regarding the state reported by a Workload Manager. There are a few states reported by a Workload Manager:
  1. Running - Workload Manager is executing as expected.
  2. Hibernating - Workload Manager is operating as a backup scheduler
  3. Unknown - Contact with Workload Manager has failed or has not been correctly set up.
- Host - The host name where Workload Manager is located.
- Port - The port which the Workload Manager communicates.

3.3. Nodes

3.3.1. Create Node Profile

Summary

A *node* is a computer consisting of 1 or more processors. A *node profile* is the additional information Workload Manager tracks about a specific node on the cluster. All nodes found by Workload Manager will automatically have a node profile created for them, so node profiles only need to be created for nodes that Workload Manager doesn't know exist. For example, if a system administrator were planning on adding 32 new nodes to the cluster, the system administrator could create all 32 node profiles before the nodes were added to the cluster.

Create Node Information		
Field Name	***Required***	***Description***
Node ID	Required	This field assigns the node profile to a desired node.
Node Speed	Optional	This field allows a user the option of specifying the relative speed of this node in comparison to other nodes. By default a value of 1.0 is given to all the nodes on the cluster. If a subset of nodes are faster than the the rest of the cluster, a higher speed should be given to them. The *node speed* values are determined by the system administrator and are not based upon any information gathered by Workload Manager.

Create Node Information		
Processor Speed	Optional	This field allows a user the option of specifying the processor(s) speed on this node. This provides Workload Manager the information needed to schedule nodes with similar processor speeds.
Rack Number	Optional	This field allows a user the option of specifying the rack number where the node is located.
Slot Number	Optional	This field allows a user the option of specifying the slot number where the node is located.
Class / Queue	Optional	This field allows a user the option of specifying the classes/queues that can access this node.
Partitions	Optional	This field allows a user the option of specifying partitions to which a node is assigned.
Features	Optional	This field allows a user the option of specifying features assigned to this node.

Maximum Node Limits		
Field Name	***Required***	***Description***
Maximum Jobs on Node	Optional	This field allows the user the option of specifying the maximum number of simultaneous jobs allowed to run on this node.
Maximum Jobs Per User on Node	Optional	This field allows the user the option of specifying the maximum number of simultaneous jobs per end user allowed to run on this node.
Maximum Load on Node	Optional	This field allows the user the option of specifying the maximum percentage of load allowed to run on this node. *Load* is the number of jobs divided by the number of processors.

3.3.2. Modify A Node Profile

Summary

A *node* is a computer consisting of one or more processors. A *node profile* is the additional information Workload Manager tracks about a specific node on the cluster. This window allows a user to modify the node profile.

Node Information		
Field Name	***Displayed***	***Description***
Node ID	Always	This field assigns the node profile to a desired node.

General Attributes		
Field Name	***Displayed***	***Description***
Node Speed	Data Dependent	This field allows a user the option of specifying the relative speed of this node in comparison to other nodes. By default a value of 1.0 is given to all the nodes on the cluster. If a subset of nodes are faster than the the rest of the cluster, a higher speed should be given to them. The node speed values are determined by the system administrator and are not based upon any information gathered by Workload Manager.
Processor Speed	Data Dependent	This field allows a user the option of specifying the processor(s) speed on this node. This provides Workload Manager the information needed to schedule nodes with similar processor speeds.
Partition	Data Dependent	This field allows a user the option of specifying partitions to which a node is assigned.

Node Usage Limits		
Field Name	***Displayed***	***Description***

Node Usage Limits		
Maximum Jobs on Node	Always	This field allows the user to specify the maximum number of simultaneous jobs allowed to run on this node.
Maximum Jobs Per User on Node	Always	This field allows the user to specify the maximum number of simultaneous jobs per end user allowed to run on this node.
Maximum Load on Node	Always	This field allows the user the option of specifying the maximum percentage of load allowed to run on this node. *Load* is the number of jobs divided by the number of processors.

Cluster Summary		
Field Name	***Displayed***	***Description***
Available Class	Data Dependent	This field allows a user the option of specifying the classes that can access this node.
Replace / Append Features	Data Dependent	This field allows a user the option of specifying features assigned to this node.

3.3.3. List Nodes

Categories

Summary:Node ID, State, Class, Features, Job List, Messages, Operating System List, Total Processors

Description:Node ID, State, Features, Network, Node Type, Operating System, Operating System List, Partition, Processor Speed, Rack, Size, Slot, Speed

Configured Resources:Node ID, State, Total Disk, Total Memory, Total Processors, Total Swap

Available Resources:Node ID, State, Available Disk, Available Memory, Available Processors, Available Swap

Usage Limits:Node ID, State, Maximum Input/Output In, Maximum Input/Output Load, Maximum Input/Output Out, Maximum Jobs, Maximum Jobs Per User, Max Load, Maximum Processor Equivalent

Per Job, Maximum Processors, Maximum Processors Per Class

Diagnostics:Node ID, State, Load, Messages, Reservation Count, Block Reason

***Comments*:** Node ID, State, Comments

List Nodes Fields		
Field	***Categor ies***	***Additional Information***
Node ID	All	All nodes require a unique ID. This field displays that ID.
State	All	This field displays the operating status of the node. For example, unknown, draining, busy, running, down, idle, etc.
Architecture		This field displays the hardware architecture of the node. The exact hardware information displayed will depend upon the information the resource manager supplies to Workload Manager.
Consumable Resources		This field displays the restricted resources and the current number available for the node. Some clusters have restrictions placed upon certain nodes. Usually these restrictions are in the form of software licenses. Sometimes a software license can restrict the number of jobs that can simultaneously be using the software on a node. *Consumable resources* allow a system administrator to define the number of licenses or other restricted resources available on a particular node.
Available Class		This field displays the classes that can access the node.
Class	Summary	This field displays the classes that can access the node.

List Nodes Fields		
Features	Summary, Description	A feature is a custom attribute often describing a unique hardware or software configuration associated with the node. This field displays the features associated with the node.
Frames		This field displays the rack / frame number where the node is logically located. This field is only availably for backward compatibility with older versions of Workload Manager. Refer to the rack field for this information.
Job List	Summary	A node can execute one or more jobs simultaneously. This field displays a list of jobs currently executing on the node.
Load	Diagnostics	The *load* is defined as the number of processors on the node divided by the number of jobs on the node. This field displays what the current node load is.
Maximum I.O. Input	Usage Limits	This field displays the maximum disk input in bytes that can occur before the node state is changed to busy.
Maximum I.O. Load	Usage Limits	This field displays the maximum disk input and output in bytes that can occur before the node state is changed to busy.
Maximum I.O. Output	Usage Limits	This field displays the maximum disk output in bytes that can occur before the node state is changed to busy.
Maximum Jobs	Usage Limits	This field displays the maximum number of jobs allowed on the node at one time.
Maximum Jobs Per User	Usage Limits	This field displays the maximum number of jobs for a single user allowed on the node at one time.

List Nodes Fields		
Maximum Load	Usage Limits	The *load* is defined as the number of processors on the node divided by the number of jobs on the node. This field displays the maximum load for the node.
Maximum Processor Equivalent Per Job	Usage Limits	This field displays the maximum number of processor equivalents* per job allowed on this node at one time.
Maximum Processors	Usage Limits	This field displays the maximum number of utilized processors allowed on this node at one time.
Maximum Processors Per Class	Usage Limits	This field displays the maximum number of utilized processors per class allowed on this node at one time.
Messages	Diagnostics	This field displays information messages provided by Workload Manager relating to the node.
Network	Description	This field displays the network hardware on the node.
Node Type	Description	A *node type* is a custom tag attached to a node. It is usually used in conjunction with an allocation manager such as QBank to assign different charge rates according to the specific node type. This field displays the node type attached to the node.
Operating System	Description	A node is configured with a specific operating system. This field displays the node's configured operating system.
Operating System List	Summary, Description	A node is configured with a specific operating system. This field displays the node's configured operating system, as well as other operating systems that are compatible with the configured operating system.

List Nodes Fields		
Partition	Description	Clusters can be divided into different sections. These sections are commonly called *partitions*. This field displays the partition to which the node is assigned.
Priority Function	Description	This field displays which priority function will be used to calculate a node's priority.
Priority		This field displays the priority of the node. The default priority is 0.
Processor Speed		This field displays the processor speed as gathered from the resource manager.
Rack	Description	This field displays the rack number where the node is logically located.
Available Disk	Available Resources	This field displays the available disk space, measured in megabytes (MB), on the node.
Available Memory	Available Resources	This field displays the available memory, measured in megabytes (MB), on the node.
Available Processors	Available Resources	This field displays the number of available processors on the node.
Available Swap	Available Resources	This field displays the available swap, measured in megabytes (MB), on the node.
Total Disk	Configured Resources	This field displays the total disk space, measured in megabytes (MB), on the node.
Total Memory	Configured Resources	This field displays the total memory, measured in megabytes (MB), on the node.
Total Processors	Summary, Configured Resources	This field displays the total number of processors on the node.
Total Swap	Configured Resources	This field displays the total swap space, measured in megabytes (MB), on the node.
Reservation Count	Diagnostics	This field displays the number of reservations on the node.

List Nodes Fields		
Size	Description	This field displays a description of the size of the node such as 1u or 2u.
Slot	Description	This field displays the slot number where the node is logically located.
Node Speed	Description	This field displays how much faster this node is from the default 1.0 node. For example, if this node were 50% faster than the default node this field would display 1.5. The *node speed* is used to determine proper wall clock limits and CPU time scaling adjustments.
Statistics Active Time		This field displays the total time the node has actively been executing jobs.
Statistics Total Time		This field displays the total time the node has been on the cluster.
Statistics Up Time		This field displays the total time the node has been available to execute jobs.
Block Reason	Diagnostics	This field displays any error messages related to the node.
Comments	Comments	This field gives the user the option of attaching a comment to the node.

3.3.4. Node Calendar

Summary

The node calendar displays the jobs and reservations on a calendar. The top bar or x-axis is the displayed time frame. The left bar or y-axis is the nodes on the cluster. The colored boxes/cells in the table are identified in the display key.

The node calendar supports 4 time frames. Days in Month, Days in Week, Hours in Day, Minutes In Hour. The top left tabs allow the user to choose the desired time frame. When the *display selected time frame* button is selected, the desired time frame will be displayed.

The Display Key panel allows the user to show or hide resources on the node calendar. The resize calendar panel, when enabled, reduces the size of each box/cell in the table to allow the user to see more nodes at once. The node names are not visible when the table is compacted.

The Current Cluster Time allows a user to see what the current time on the cluster is.

3.3.5. Node Timeline

he Node Timeline window displays the jobs and reservations executing on each node. On the left side of the Now line is the amount of time completed for the jobs and reservations, while the right side is the remaining execution/reservation time. When the cursor arrow is held over the timelines, Cluster Manager displays the job ID number.

3.4. Partitions

3.4.1. Create a Partition Profile

Clusters can be divided into different sections. These sections are commonly called *partitions*. A partition is a semi-permanent division of the cluster and is most often created for certain nodes containing unique hardware. It should be recognized that jobs are not allowed to run in more then one partition. If jobs need the ability to span multiple partitions an infinite reservation should be used instead of a partition.

Partition Information		
Field	**Required**	**Additional Information**
Partition Name	Required	This field allows a user the ability to create a custom name for this partition.
User	Optional	This field is used to specify which users can access this partition.
Group	Optional	This field is used to specify which groups can access this partition.
Class/Queue	Optional	This field is used to specify which classes /queues can access this partition.
Account	Optional	This field is used to specify which accounts can access this partition.

Partition Information		
Quality of Service (QoS)	Optional	This field is used to specify which quality of services *(*QoS *)* can access this partition.
Node	Optional	This field is used to specify which nodes are members of this partition.

3.4.2. Modify a Partition Profile

Clusters can be divided into different sections. These sections are commonly called *partitions*. A *partition* is a semi-permanent division of the cluster and is most often used when certain nodes contain unique hardware. It should be recognized that jobs are not allowed to run in more than one partition. If jobs need the ability to span multiple partitions an infinite reservation should be used instead of a partition.

The ability to modify partitions is not currently available in Cluster Manager.

3.4.3. List Partitions

Summary

This tool lets you view additional information about partitions. To learn more about what partitions are and how they are created, please see the Documentation.

Default Display Categories

Summary:Partition ID, Resource Manager

Disk:Partition ID, Total Disk, Reserved Disk, Utilized Disk

Memory:Partition ID, Total Memory, Reserved Memory, Utilized Memory

Swap:Partition ID, Total Swap, Reserved Swap, Utilized Swap

Nodes:Partition ID, Node List, Total Nodes, Reserved Nodes, Utilized Nodes

Processor:Partition ID, Total Processors, Reserved Processors, Utilized Processors

Credentials:Partition ID, User Access List, Group Access List, Account Access List, Class Access List, Quality of Service(QoS) Access List

List Partition Fields		
Field	**Categories**	Additional Information
Partition ID	All	This field displays the partition ID, or the name of the partition.
Resource Manager	Summary	This field displays the resource manager ID of which this partition is a member.
Total Disk	Disk	This field displays the total disk space available in the partition. The disk space is measured in megabytes (MB).
Reserved Disk	Disk	This field displays the amount of disk space reserved by this partition. The disk space is measured in megabytes (MB).
Utilized Disk	Disk	This field displays the amount of disk space currently being used by this partition. The disk space is measured in megabytes (MB).
Total Memory	Memory	This field displays the total memory available in this partition. The memory is measured in megabytes (MB).
Reserved Memory	Memory	This field displays the amount of memory reserved in this partition. The memory is measured in megabytes (MB).
Utilized Memory	Memory	This field displays the amount of memory currently being used by this partition. The memory is measured in megabytes (MB).
Total Swap Space	Swap	This field displays the total swap space available in the partition. The swap space is measured in megabytes (MB).
Reserved Swap	Swap	This field displays the amount of swap space reserved by this partition. The swap space is measured in megabytes (MB).

List Partition Fields		
Utilized Swap	Swap	This field displays the amount of swap space currently being used by this partition. The swap space is measured in megabytes (MB).
Node List	Nodes	This field displays the names of the nodes available in this partition.
Total Nodes	Nodes	This field displays the total number of nodes available in this partition.
Reserved Nodes	Nodes	This field displays the number of nodes reserved in this partition.
Utilized Nodes	Node	This field displays the number of nodes currently being used in this partition.
Total Processors	Processor	This field displays the total number of processors available in this partition.
Reserved Processor	Processor	This field displays the number of processors reserved in this partition.
Utilized Processor	Processor	This field displays the number of processors currently being used in this partition.
User Access List	Credentials	This field displays the users that can access this partition.
Group Access List	Credentials	This field displays the groups that can access this partition.
Account Access List	Credentials	This field displays the accounts that can access this partition.
Class Access List	Credentials	This field displays the classes that can access this partition.
Quality of Service(QoS) Access List	Credentials	This field displays the quality of services that can access this partition.

3.5. Licenses

3.5.1. List Licenses

Licenses are reported to Moab via a license manager such as FlexLM. Each license is treated as a generic resource that can be consumed if specified in job submission. This license table is meant to help maximize license usage and provide useful information regarding licenses. Below is a table of terms explaining each field found in the license list.

Default Categories

Summary:License, Available, Configured

History:License, Idle %, Busy %, Avg. In Use

List Licenses Fields

Field	*Category*	*Field Information*
License	All	All licenses have a unique name specified as a generic resource from a license manager. This name is displayed here.
Available	Summary	This displays the number of licenses that are currently free to use.
Configured	Summary	This displays the number of licenses configured for use for each individual license type.
Idle %	History	This is calculated using RM poll intervals. This takes each interval where none of the licenses are being used divided by total RM intervals. (free iterations / total iterations * 100)
Busy %	History	This is calculated using RM poll intervals. This takes each interval where all of the licenses are being used divided by total RM intervals. (busy iterations / total iterations * 100)

List Licenses Fields		
Avg. In Use	History	This is calculated using RM poll intervals. This takes the total number of licenses being used per iteration divided by total RM intervals. (total licenses / total iteration)

In addition to a list of license information, it is also possible to display historical statistical information. There are three types of statistical displays on the bottom left hand side of the window:

- **License State Percentage** - A pie chart is listed for each license. In the chart, idle (no licenses used), active (some licenses used) and busy (all licenses used) iterations are displayed.
- **Total Usage Ratio** - A bar chart with every bar representing a license. The usage ratio is (used licenses / total licenses).
- **Usage Ratio Over Time** - A line graph with lines representing each license RM. The usage ratio is (used licenses / total licenses).

3.6. Cluster

3.6.1. Visual Cluster

Summary

The visual cluster gives an easy and concise way of viewing your entire cluster and the status of each node. The table and explanation below explain how to interpret the visual cluster:

Table 3-1. Visual Cluster Example

	Slot 1	Slot 2	Slot 3
Rack 1	Node A		Node D
Rack 2		Node B	
Rack 3			Node C

A rack is a physical frame that holds a node. The slot is the location of the node inside the rack. The racks make up the first column of the table. The slot locations increase from left to right. For example, Node A is located on Rack 1 in Slot 1. Node D is also located on Rack 1 but instead of Slot 1 it's located in Slot 3. In the visual cluster Node A through Node D are displayed as icons. The different icons can represent node state, node attributes, reservations, jobs, and/or nodes. The subpanel sections below

describe these states in more detail. Further information can be gathered about nodes by hovering the mouse over any nodes.

It should be noted that the visual cluster is for display purposes only and the location of the node does not play any part in how Workload Manager schedules.

3.6.1.1. Node Attribute Selector

The node attribute selector gives the user the power to see various attributes of the nodes displayed in the Visual Cluster. This allows the user to compare and contrast attributes of interest. Node attributes include standard categories such as architecture, OS, hardware metrics (memory, disk, swap, etc.), as well as any metric read in through Moab as a generic metric (a node's GMETRIC). The default display for node attributes is the corresponding color of the outer rim of each node cell. This can be changed in the section titled "Node Display Options".

The "Clear Attribute" button will simply clear any selection and node attribute displayed currently. The "Graph Attributes" button will display each node attribute based on state, current load, or historical load. This is tied to the value currently selected in the "Node Usage Display Options".

Once a node attribute is selected, Moab Cluster Manager will determine the number of nodes and processors that describe each attribute and will display a corresponding key value that matches the Visual Cluster by color. Each attribute's display can be individually controlled via the check box next to each node attribute name and color.

If the node attribute is a numerical value - for example, a generic metric - then Moab Cluster Manager will attempt to place the values into a reasonable range as to effectively categorize the values.

3.6.1.2. Highlight Reservations, Jobs, and/or Nodes

This section provides a user with the option of highlighting resources in the visual cluster table. The three sections are divided into reservations, jobs, and nodes. Each section can be simultaneously displayed by having each border being a different color. The white box displays the names of the selected resources with the headers "Res.: ", "Job: ", and "Node: " respectively. The Select buttons open the lists of reservations, list jobs, or list nodes window depending on which resource the user has selected. The visual cluster window will appear with the desired resources highlighted. The Clear buttons remove the highlight from the visual cluster table and erases the names from the colored box. The Color button changes the highlight color for the specified resource. The new highlight color will be displayed in the colored box.

3.6.1.3. Node Display Options

There are different options depending on how the three checkboxes - "Hide Usage", "Hide Attributes",

and "Auto Resize" - are set:

- Usage unchecked, attributes unchecked (DEFAULT) - Usage will be displayed on the inside of the cell and attributes will be displayed on the outside border of the cell.
- Usage checked, attributes unchecked - Usage will not be displayed, and attributes will take up the entire node cell.
- Usage unchecked, attributes checked - Usage will take up the entire node cell, and attributes will not be displayed.
- Usage checked, attributes checked - No information will be displayed, leaving each cell grayed out.
- Auto resize checked - The table of nodes will try to fix to the size of the window given. If there are more nodes than can fit on the window, a minimum size will be set with the rest of the nodes dangling off the window.
- Auto resize unchecked - The table of nodes will always be set to a minimum size. However, a horizontal scroll bar will be set if the nodes dangle off the side of the window.

The "Node Height" slider changes the height of each row to grow or shrink the nodes to fit the user's display needs.

Highlight reservations, jobs, and nodes remain unaffected regardless of what node display options are set.

3.6.1.4. File Menu Options

Note: This menu is also accessible by right clicking anywhere on the main window.

Actions Menu Options

- Online Node - This option will change a node's status from unavailable to available. An online node is available for jobs to execute on it.
- Offline Node - This option will change a node's status from available to unavailable. An offline node is unavailable for jobs.
- Reserve Selected Nodes - This option will prepopulate the desired nodes in a create reservation window with the nodes that were selected using the mouse.
- Reserve Highlighted Nodes - This option will prepopulate the desired nodes in a create reservation window with the nodes that were selected using the Node Attribute Selector.
- Modify Nodes - This option will open a modify node(s) window that will allow the administrator to modify one selected node or perform group operations over numerous selected nodes.

- Power On Nodes - This option will change the power status of the selected nodes to ON. To take advantage of this command, CLUSTERQUERYURL and NODEPOWERURL must be set up to handle xCAT or IPMI interfaces.
- Power Off Nodes - This option will change the power status of the selected nodes to OFF. To take advantage of this command, CLUSTERQUERYURL and NODEPOWERURL must be set up to handle xCAT or IPMI interfaces.
- Reboot Nodes - This option will change the power status of the selected nodes to REBOOT. To take advantage of this command, CLUSTERQUERYURL and NODEPOWERURL must be set up to handle xCAT or IPMI interfaces.

Highlight Menu Options

- Highlight Jobs for Selected Nodes - This will get the name of the jobs from the selected node and highlight all the nodes that that job is on.
- Highlight Reservations for Selected Nodes - This will get the name of the reservations from the selected node and highlight all the nodes the reservation is on.
- Select Nodes with Credential - This will select each node running a job with the selected credential.

Display Menu Options

- View Processor Usage - This will open a new window that displays processor usage.

3.6.1.5. Workload Manager Usage Break Down

The default display for usage breakdown is the inner core of the node cell. This can be changed in the section Node Display Options.

There are three options for displaying usage breakdown:

- Display Node State
- Display Current Load
- Display Historical Load

3.6.1.5.1. Display Node State

Display node state displays the state the node is in according to the Workload Manager.

- Down - The node is currently reporting a state of "Down" because of failure or administrative action.
- Full Load - The node is currently reporting a state of "Busy".
- Partial Load - The node is currently reporting a state of "Running".
- Unused - This is currently unused by node state.
- Offline - The node is currently reporting a state of "Offline." This is also the default sate when the state is not recognized.
- Idle - The node is currently reporting a state of "Idle".

3.6.1.5.2. Display Current Load

Display historical load displays the average percentage over time that the node has been used.

- > 100 % - The node is currently executing more executables than it has processors
- 80% - 100% - The node is currently executing executables on between 80 and 100 percent of its processors.
- 60% - 80% - The node is currently executing executables on between 60 and 80 percent of its processors.
- 40% - 60% - The node is currently executing executables on between 40 and 60 percent of its processors.
- 20% - 40% - The node is currently executing executables on between 20 and 40 percent of its processors.
- 0% - 20% - The node is currently executing executables on between 0 and 20 percent of its processors.

3.6.1.5.3. Display Historical Load

Display historical load displays the average percentage over time that the node has been used.

- > 100 % - The node has historically executed more executables than it has processors
- 80% - 100% - The node has historically executed executables on between 80 and 100 percent of its processors.
- 60% - 80% - The node has historically executed executables on between 60 and 80 percent of its processors.
- 40% - 60% - The node has historically executed executables on between 40 and 60 percent of its processors.
- 20% - 40% - The node has historically executed executables on between 20 and 40 percent of its processors.

- 0% - 20% - The node has historically executed executables on between 0 and 20 percent of its processors.

3.6.2. Processor Usage

Summary

This graph displays how the cluster's processors are being used over time. The left bar, or y-axis, displays the number of processors. The bottom bar, or x-axis, displays time. The light yellow color displays the total available processors on the cluster. The dark yellow color displays the processors used by jobs and job reservations. The blue color displays the processors used by reservations other than job reservations.

The switch statstics option allows for "Available Processors" and "Jobs Reservations" colors to be switched.

3.6.3. Add Resource Manager

Summary

As the name suggests a resource manager manages compute resources. Different resource managers manage different resources. Possible resources are hardware, software licenses, storage, networks, or compute cycles.

3.6.3.1. Resource Manager Add Options

- Resource Manager Type - This field displays the type of resource manager interface that is being enabled.
- Name - This field displays the unique resource manager name.
- Description - This field displays a description of what the resource manager does.
- Port - This field allows an administrator to select the port on which Workload Manager will communicate with this resource manager.
- Server URL - This field allows an administrator to input the URL of the resource manager. A URL must be entered in one of the following formats:

  - File://[File Path] - This field requires a file that acts as a resource manager. For example, if a file called rmfile.txt were located in the tmp directory, then the format would be File://tmp/rmfile.txt

- http://[address] = This field requires the web address of the resource manager. For example, if the resource manager were located at 10.10.10.100 then the format would be http://10.10.10.100
- [PATH]/executable This field requires an executable. For example, if the resource manager were rm.sh, located in the tmp directory, then the format would be /tmp/rm.sh

3.6.4. List / Modify Resource Managers

Summary

As the name suggests a resource manager manages compute resources. Different resource managers manage different resources. Possible resources are hardware, software licenses, storage, networks, or compute cycles.

3.6.4.1. List Resource Manager Fields

- Resource Manager Name - This field displays the custom name given to the resource manager by the system administrator.
- Resource Manager Type - This field displays the type of resource manager interface enabled.
- Resource Manager State - This field displays the status of the resource manager. Possible states include active, idle, ordown.

3.6.4.2. Modify Resource Manager

- Resource Manager Type - This field allows an adminstrator to change the resource manager interface.
- Server URL - This field allows an administrator to input the URL of the resource manager. A URL must be entered in one of the following formats:

  - File://[File Path] - This field requires a file that acts as a resource manager. For example, if a file called rmfile.txt were located in the tmp directory, then the format would be File://tmp/rmfile.txt
  - http://[address] = This field requires the web address of the resource manager. For example, if the resource manager were located at 10.10.10.100 then the format would be http://10.10.10.100
  - [PATH]/executable This field requires an executable. For example, if the resource manager were rm.sh, located in the tmp directory, then the format would be /tmp/rm.sh

- Name - This field allows an administrator to change the current resource manager name given to this resource manager interface.
- Port - This field allows an administrator to select the port on which Workload Manager will communicate with this resource manager.
- State - This field displays the current state of the resource manager interface.
- Total Requests - This field displays the total number of communications that have occurred between Workload Manager and the resource manager.
- Response Time (In Seconds) - This bar graph displays the average response time, as well as the maximum response time between Workload Manager and the resource manager. This information often provides valuable diagnostic information when resource manager errors are occurring.

3.6.5. Resource Manager Messages

Summary

Resource managers have the ability to report diagnostic messages and user specified messages. These messages can be used to gain further information or knowledge about a particular resource manager. This may be useful in trying to diagnose failures associated with the resource manager.

Resource managers' messages are divided into three categories: a diagnostic message, other messages, and peer service interface messages. All message types are described in greater detail below.

3.6.5.1. Resource Manager Diagnostic Message

The first field in the resource manager messages frame is the diagnostic message. This diagnostic message reports any problems that Moab may see with the resource manager configuration. Examples include missing resource manager parameters or parameters that are malformed.

3.6.5.2. Resource Manager General Messages

The second field is table of messages attached to the resource manager itself. These messages may be user specified messages that describe notes about the resource manager. They may also be generalized system messages Moab generates that summarize issues going on with the resource manager itself. The order that messages are appear are from oldest to newest.

3.6.5.3. Resource Manager Peer Service Interface (PSI) Messages

The third field is also a table of messages, but it reports very specific information concerning the resource manager's peer service interface. This is the module inside the resource manager that is responsible for communicating with Moab and other resource managers. PSI messages consist of three parts:

- Type - This is the type of failure reported by the message. Some types include "clusterquery", "workloadquery", or "rminitialize".
- Time - This is the reported time of the message.
- Message - This is the actual messsage text itself.

3.6.6. Allocation Manager

Summary

An allocation manager functions much like a bank in that it provides a form of currency which allows jobs to run on a cluster. Each job on the cluster requires a certain number of credits to be eligible to execute. An allocation manager tracks the used credits and notifies Workload Manager of any jobs that would exceed their credit limit.

3.6.6.1. External Allocation Manager Settings

- Name - This field allows an administrator to define a name for the Allocation Manager.
- Hostname - This field allows an administrator to input the URL of the resource manager. A URL must be entered in one of the following formats:

  - File://[File Path] - This field requires a file that acts as a resource manager. For example, if a file called rmfile.txt were located in the tmp directory, then the format would be File://tmp/rmfile.txt
  - http://[address] = This field requires the web address of the resource manager. For example, if the resource manager were located at 10.10.10.100 then the format would be http://10.10.10.100
  - [PATH]/executable This field requires an executable. For example, if the resource manager were rm.sh, located in the tmp directory, then the format would be /tmp/rm.sh

- Port - This field allows an administrator to select the port on which Workload Manager will communicate with this allocation manager.
- Timeout - This field allows an administrator to define how long Workload Manager will wait for the Allocation Manager to respond to messages.

- Type - This field allows an administrator to define which allocation manager type is being used. The following options are available.

    1. Gold
    2. GGF
    3. Qbank
    4. ResD
    5. File

- Allocation Failure Job Action - This field allows an administrator to define what should happen to a job if an allocation manager failure is detected. The following options are available.

    1. Log Failure
    2. Reattempt

- Wire Protocol - This field allows an administrator to define which wire protocol will be used by Workload Manager to communicate with the Allocation Manager. The following options are available.

    1. Default
    2. HTML
    3. XML
    4. SSS2

- Socket Protocol - This field allows an administrator to define which socket protocol will be used by Workload Manager to communicate with the Allocation Manager. The following options are available.

    1. HTTP
    2. SSS-HALF
    3. SSS-Challenge

- Secret Key - This field allows an administrator to encrypt communication between the allocation manager and Cluster Manager using a secret key.

- Append Machine Name - If this field is enabled, Cluster Manager will append the machine name to each account before submitting debits to the allocation manager. This will create unique charges per machine name.
- Charge Rate Policy - This field allows an administrator to define how charging per job occurs. The following options are available.

    1. DebitAllWC - This option will debit from the allocation manager according to the time used on the cluster.
    2. DebitAllCPU - This option will debit from the allocation manager according to how many processors are used and for how long the processors are used.
    3. DebitAllPE - This option will debit from the allocation manager according to processor equivalent [1] seconds.
    4. DebitSuccessfulWC - This option will debit from the allocation manager when a job successfully completes execution according to the amount of time used on the cluster.
    5. DebitSuccussfulCPU - This option will debit from the allocation manager when a job successfully completes execution according to how many processors are used and for how long the processors are used.
    6. DebitAllPE - This option will debit from the allocation manager when a job successfully completes execution according to processor equivalent seconds.

- Flush Interval - This field allows an administrator to define how long Workload Manager will wait before contacting the allocation manager.
- Fall Back Account - This field allows an administrator to define a second account jobs can use if their allocation manager account doesn't have adequate resources to allow the job to start executing. If the second account isn't defined or doesn't have adequate resources, the job is then placed on hold.

3.6.6.2. Internal Allocation Manager Settings

- Assign / Modify Fixed Allocations - This field opens the List Credentials window where throttling policies can be set. The throttling policies can be used to create fixed or unchanging restrictions on a credential.
- Assign / Modify Rolling Allocations - This field opens the Fairshare window where fairshare targets can be set. The fairshare window can be used to create rolling or interval-based restrictions on a credential.

3.7. Grid

3.7.1. Grid Summary

Note: This feature is exclusive to Moab Grid Manager. Moab Cluster Manager does not display this feature.

Summary

Grid summary displays all clusters that can be seen by the Moab the user is associated with. Helpful cluster information is displayed as well

Default Categories

Summary: Cluster Name, Host, Port, Average Response Time, Relationship, State

Authorization: Cluster Name, Authorization

Resources: Cluster Name, Idle Nodes, Total Nodes, Idle Processors, Total Processors

Credentials: Cluster Name, Class, Account, QoS

Cluster Profile: Cluster Name, Architecture, O.S., Network Type, Node Features

Flags: Cluster Name, Rsv. Export, Rsv. Import, Collapsed Node View, Local Workload Export

Grid Summary Fields		
Field	***Category***	***Field Information***
Cluster Name	ALL	This field displays the remote cluster's name. The cluster name is determined by the SCHEDCFG parameter's name.
Host	Summary	The name of the host where the remote Moab is located.
Port	Summary	The port of the host where the remote Moab is located.

Grid Summary Fields		
Average Response Time	Summary	This field displays the average time it takes for the cluster currently logged into to communicate with the listed cluster.
Relationship	Summary	The relationship the remote cluster has to the local resources.
State	Summary	The state of the remote cluster according to this cluster. "Active" means the connection is healthy; "Corrupt" means the connection configuration is incorrect or another connection problem has occurred.
Authorization	Authorization	The level of control/information available to requests coming from this source/peer.
Idle Nodes	Resources	This field displays the number of nodes on the listed cluster not being used.
Total Nodes	Resources	This field displays the total number of nodes on the listed cluster.
Idle Processors	Resources	This field displays the number of processors not being used on the listed cluster.
Total Processors	Resources	This field displays the total number of processors on the listed cluster.
Architecture	Cluster Profile	This field lists all the node architectures detected on the listed cluster. The architecture of a node can be specified via the NODECFG parameter.
Operating System	Cluster Profile	This field displays operating systems detected on the listed cluster. The operating system of a node can be specified via the NODECFG parameter.

Grid Summary Fields		
Network Type	Cluster Profile	This field displays the hardware network types detected on the listed cluster. The network type of a node can be specified via the NODECFG parameter.
Node Features	Cluster Profile	This field displays all node features detected on the listed cluster. Node features can be specified via the NODECFG parameter.
Class	Credentials	This field displays all classes on the listed cluster.
Account	Credentials	This field displays all accounts on the listed cluster.
QoS	Credentials	This field displays all QoA accounts on the listed cluster.
Reservation Export	Flags	Allows local reservations to be exported. The local reservations must be explicitly imported by remote clusters for them to be seen and used.
Reservation Import	Flags	Allows remote reservations to be imported. The remote reservations must be explicitly exported by remote clusters for them to be seen and used.
Collapsed Node View	Flags	The remote cluster's nodes will be collapsed into one SMP-like node locally.
Local Workload Export	Flags	The local workload will be visible to remote clusters.

3.7.2. Visual Grid

Note: This feature is exclusive to Moab Grid Manager. Moab Cluster Manager does not display this feature.

Summary

The Visual Grid is a graph showing the relationships of clusters that the user can see. The cluster currently connected to is shown in yellow in the center of the graph. Each neighboring box corresponds

to a cluster that is connected to the central cluster. There are three types of relationships between clusters:

- Peer - Represented with arrows pointing both directions; the neighbor is colored red. Workload can be directed in both directions.
- Master - Represented with an arrow pointing from the neighbor to the center cluster; the neighbor is colored blue. Workload can only be submitted to the neighbor who can schedule jobs on the central cluster.
- Slave - Represented with an arrow pointing from the center to the neighbor cluster; the neighbor is colored green. Workload can only be submitted to the center who can schedule jobs on the neighbor cluster.

Graph Features

The graph's cells and edges can be moved, so if there any arrows or cells that are blocked, feel free to move them. Zooming is also allowed via the mouse wheel. Scrolling up will zoom in, and scrolling down will zoom out. The edges can be modified via right clicking to create a new pivot. Click and drag, shift-click, and ctrl-click are all functional as well.

Visual Cluster

The Visual Cluster is accessible from the Visual Grid window. Each individual cluster's nodes can be accessed by double clicking the corresponding cluster box; this can also be done by right-clicking on the cluster box and selecting "View in Visual Cluster". If there are nodes corresponding to the cluster name, an empty Visual Cluster will be displayed. Also, the Visual Cluster will retain slot and rack information gathered from Moab Workload Manager. If you do not want to filter out any nodes, select the "Display All Nodes" button at the bottom of the window, and all nodes will be seen.

Modifying Grid Relationship

The relationships that a local cluster has to remote clusters can be modified and deleted by an administrator via the Visual Grid window. To do this, right-clicking on the remote cluster to modify will give the option of modifying or deleting the relationship. Modifying the relationship will bring up a new window titled, "Modify Grid Relationship". Deleting a relationship will remove the pertinent lines from the moab.cfg file to detach the local cluster's connection from the selected remote cluster.

View Grid Diagnostic Messages

The relationships with remote clusters may have issues from time to time for whatever reason. From the visual grid, it is possible to view these messages reported on a per cluster basis. If messages exist for a particular connection to a remote cluster, the remote cluster in question will have a warning icon with an exclamation point (!). Right-clicking on the remote cluster and selecting "View Cluster's Messages" will bring up the messages reported through the resource manager interface as seen in the "Resource Manager Messages" table.

3.7.3. Create Grid Relationship

Note: This feature is exclusive to Moab Grid Manager. Moab Cluster Manager does not display this feature.

Summary

Create Grid Relationship allows a user with level 1 Moab Admin privledges to create a connection between the current cluster and a remote cluster specified. Configuration must be done on both clusters to make the relationship valid.

Remote Cluster Information		
Field	***Required***	***Field Information***
Relationship	Yes	This field allows the local cluster to be set up according to the ***remote*** cluster's relationship to it. The types of cluster relationships are as follows: • Peer - This allows the local cluster to get information from a remote cluster without giving up control. To set up bidirectional job flow, peer should be set on both clusters. • Slave - This allows the local cluster to send information to and get control of a remote cluster. • Master - This allows the remote cluster to take control of the local cluster's resources. Jobs can only be submitted via the master node.
Scheduler Name	Yes	This field is where the remote cluster's name should be entered.

Remote Cluster Information		
Host Name	Peer or Slave Only	The IP address or host name (if known) of the remote cluster should be entered here. If the local cluster is a slave, this information is not needed.
Port Number	Peer or Slave Only	The port number of the remote cluster should be entered here. If the local cluster is a slave, this information is not needed.
Key	Yes	In order to validate a connection, the local and remote clusters need to share a private key. The key should be entered in this field on both clusters.

Grid Data Staging		
Field	***Required***	***Field Information***
Enable Data Staging	Optional	Allows grid data staging to occur on the storage manager specified.
Storage Manager	Optional	The storage manager used to stage and monitor data staging files for jobs.

3.7.4. Modify Grid Relationship

Note: This feature is exclusive to Moab Grid Manager. Moab Cluster Manager does not display this feature.

Summary

Modify Grid Relationship allows a user with level 1 Moab Admin privledges to modify a connection between the current cluster and a remote cluster specified. Configuration must be done on both clusters to make the relationship valid.

Remote Cluster Information		
Field	***Required***	***Field Information***
Scheduler Name	Not modifiable	The name of the remote cluster.

Remote Cluster Information		
Relationship	Yes	This field allows the local cluster to be set up according to the ***remote*** cluster's relationship to it. The types of cluster relationships are as follows: • Peer - This allows the local cluster to get information from a remote cluster without giving up control. To set up bidirectional job flow, peer should be set on both clusters. • Slave - This allows the local cluster to send information to and get control of a remote cluster. • Master - This allows the remote cluster to take control of the local cluster's resources. Jobs can only be submitted via the master node.
Host Name	Peer or Slave Only	The IP address or host name (if known) of the remote cluster should be entered here. If the local cluster is a slave, this information is not needed.
Port Number	Peer or Slave Only	The port number of the remote cluster should be entered here. If the local cluster is a slave, this information is not needed.
Authorization	Required	"Job Grid" gives the remote cluster ADMIN1 privledges. "Control Grid" gives the remote cluster ADMIN 2 privledges. "Information Grid" gives the remote cluster ADMIN3 privledges.

Grid Data Staging		
Field	***Required***	***Field Information***

Grid Data Staging		
Enable Data Staging	Optional	Allows grid data staging to occur on the storage manager specified.
Storage Manager	Optional	The storage manager used to stage and monitor data staging files for jobs.

Flags		
Field	***Required***	***Field Information***
Reservation Export	Optional	Allows local reservations to be exported. The local reservations must be explicitly imported by remote clusters for them to be seen and used.
Reservation Import	Optional	Allows remote reservations to be imported. The remote reservations must be explicitly exported by remote clusters for them to be seen and used.
Collapsed Node View	Optional	The remote cluster's nodes will be collapsed into one SMP-like node locally.
Local Workload Export	Optional	The local workload will be visible to remote clusters.

Notes

1. Processor equivalence is a relative measure of how much of a node is taken by a job, even if only one type of node resource is requested. For example, if a job requires 1 processor and 1GB of memory, and it is running on a 4 processor node with 1GB of memory, the PE of the job is 4. All of the processors are considered to be taken because the first job is using all of the memory, which prevents any other job from running on that node.

Chapter 4. Organization

4.1. Organization Overview

The organization section allows an adminstrator to view all credentials in the system and their various roles. New credential profiles can be added as well as modified.

4.2. Visual Credential Access

Summary

This window allows a user to visually view which credentials can access which credentials. The arrows symbolize that the credential can access the other credential. There are three distinct sections displayed in the window. The first section displayed are all the credentials that can access the second section. The arrows from the first section to the second section show this. The second section contains only one credential and this credential is the selected credential. The arrows from the second section to the third section show who the selected credential can access. The first and third section may not be displayed if there are not credentials that can access or are accessed by the second section.

- Display All Credentials - This field allows a user to display all of the credentials of a specific credential type.
- Display Listed Credentials - This field allows a user to display only the listed credentials.
- Display who can access the selected credential - This field allows a user to enable/disable whether the credentials that access the selected credentials should be displayed.
- Display who the selected credential can access - This field allows a user to enable/disable whether the credentials that the selected credential can access should be displayed.

4.3. User Access

Summary

The User Access Settings page allows an administrator to assign roles to each user in the system. These settings are saved in Workload Manager and affect command-line operations, as well as permissions within Cluster Manager. The set of default roles available is:

1. *Admin1* - Users with this role are called *administrators* . They have complete control of Workload Manager.

2. *Admin2* - Users with this role are called *operators* . They have complete control of jobs, nodes, reservations, etc., but cannot modify Workload Manager configuration parameters.
3. *Admin3* - Users with this role are called *help desk personnel* . They can control various aspects of Workload Manager, but cannot modify workload.
4. *Admin4* - Users with this role are most likely trusted or experienced users. They have a subset of Admin2 permissions that is different than Admin3.
5. *Admin5* - Users with this role can only view workload and resource information.

4.4. List Credentials

Summary

This window displays all of the information regarding users, groups, accounts, classes, and qualities of service (QoS) commonly called credentials.

List Credential Fields		
Field	Category	Additional Information
Credential	All	This field displays whether the credential is a user, group, account, class, or quality of service (QoS).
Credential Identification (ID)	All	All credentials must have an identification unique to it's credential type. This field displays the credential identification.
Group	Membership	The operating system is usually responsible for the creation of groups. This field displays the groups that this particular credential ID can access.
Group Default	Default Credentials	The *default group* is the group that will be used by this credential ID's job if no group is specified.
Class	Membership	This field displays the classes that this particular credential ID can access.
Class Default	Default Credentials	The *default class* is the class that will be used by this credential ID's job if no class is specified.

List Credential Fields		
Account	Membership	This field displays the accounts that this particular credential ID can access.
Account Default	Default Credentials	The *default account* is the account that will be used by this credential ID's job if no account is specified.
Quality of Service (QoS)	Membership	This field displays the quality of services (QoS) that this particular credential ID can access.
Quality of Service (QoS) Default	Default Credentials	The *default quality of service (QoS)* is the quality of service (QoS) that will be used by this credential ID's job if no quality of service (QoS) is specified.
Partition	Partition & Reservation	Clusters can be divided into different sections. These sections are commonly called *partitions*. This field displays the partitions this credential ID can access.
Partition Default	Default Resources	The *default partition* is the partition that will be used by this credential ID's job if no partition is specified.
Credential Priority	Priority	Priority is used to decide which jobs execute first. The credential priority field allows a system administrator the ability to give certain credential IDs higher priorities over other credential IDs.
Utilized Jobs	Utilized Resources	This field displays the number of jobs currently executing for this credential ID.
Soft Maximum Job	Soft Maximum Limits	The soft maximum job limit will restrict the number of jobs allowed to execute for this credential ID. If, however, there are additional resources available after all the soft maximum job limits are met, then the hard maximum job limits are used.

List Credential Fields		
Hard Maximum Job	Hard Maximum Limits	The hard maximum job limit will restrict the number of jobs allowed to execute for this credential ID.
Maximum Job Default	Default Resources	The default maximum job is the maximum job value that will be used by this credential ID's job if no maximum job is specified.
Utilized Processors	Utilized Resources	This field displays the number of processors currently being used by this credential ID's jobs.
Soft Maximum Processors	Soft Maximum Limits	The soft maximum processor limit will restrict the number of processors used by any job for this credential ID. If, however, there are additional resources available after all the soft maximum processor limits are met, then the hard maximum processor limits are used.
Hard Maximum Processors	Hard Maximum Limits	The hard maximum processor limit will restrict the number of processors used by any job for this credential ID.
Maximum Processor Default	Default Resources	The *default maximum processors* is the maximum processors value that will be used by this credential ID's job if no maximum processor is specified.
Utilized Nodes	Utilized Resources	This field displays the number of nodes currently being used by this credential ID's jobs.
Soft Maximum Nodes	Soft Maximum Limits	The soft maximum node limit will restrict the number of nodes used by any job for this credential ID. If, however, there are additional resources available after all the soft maximum node limits are met, then the hard maximum node limits are used.
Hard Maximum Nodes	Hard Maximum Limits	The hard maximum node limit will restrict the number of nodes used by any job for this credential ID.

List Credential Fields		
Maximum Nodes Default	Default Resources	The *default maximum nodes* is the maximum nodes value that will be used by this credential ID's job if no maximum nodes value is specified.
Utilized Processor Seconds	Utilized Resources	*Utilized processor seconds* are the total number of processors used by executing jobs for this particular credential ID times the number of seconds each processor has been used.
Soft Maximum Processor Seconds	Soft Maximum Limits	The soft maximum processor seconds limit will restrict the number of processor seconds used by any job for this credential ID. If, however, there are additional resources available after all the soft maximum processor seconds limits are met, then the hard maximum processor seconds limits are used.
Hard Maximum Processor Seconds	Hard Maximum Limits	The hard maximum processor seconds limit will restrict the number of processor seconds used by any job for this credential ID.
Maximum Processor Seconds Default	Default Resources	The *default maximum processor seconds* is the maximum processor seconds value that will be used by this credential ID's job if no maximum processor seconds value is specified.
Fairshare Type	Fairshare	Refer to the fairshare section for information regarding fairshare type.
Fairshare Target	Fairshare	Refer to the fairshare section for information regarding fairshare target.
Enable Statistics	Statistics	Statistics are tracked for each credential ID. This field allows the user the option of enabling/disabling statistics for each credential ID.

List Credential Fields		
Utilized Resource Cost		What Workload Manager charges for each resource unit consumed/utilized by a job.
Dedicated Resource Cost		What Workload Manager charges for each resource unit dedicated - whether used or not - to a job.
Quality of Service (QoS) Flags		This field displays the quality of service (QoS) settings for this credential ID.
Expansion Factor Weight		This field displays the quality of service weight factor. If an idle job is submitted to this quality of service, it's expansion factor will be multiplied by this value. This will increase the job's start priority.
Queue Time Weight		This field displays the quality of service weight factor. If an idle job is submitted to this quality of service, the number of minutes that it has been in the queue will be multiplied by this value. This will increase the job's start priority.
Access Resources Queue Time Threshold		This field displays the number of minutes that an idle job must wait before it can have access to any reservations owned by the Quality of Service.
Preemption Queue Time Threshold		This field displays the number of minutes that an idle job must wait before it will be given preemptor access.
Create Reservation Queue Time Threshold		This field displays the number of minutes that an idle job must wait before a job reservation will be created for it. A job reservation will guarantee it specific resources, as well as a specific start time.

List Credential Fields		
Create Reservation Expansion Threshold		This field displays the expansion factor value that an idle job must be equal to or greater than before a job reservation will be created for it. A job reservation will guarantee it specific resources, as well as a specific start time.
Preemption Expansion Factor Threshold		This field displays the expansion factor value that an idle job must be equal to or greater than before it will be given preemptor access.
Access Resources Expansion Factor Threshold		This field displays the expansion factor value that an idle job must be equal to or greater than before it can have access to any reservations owned by the Quality of Service.
Total Credits	Credits	This field displays the total credits available to the credential ID.
Used Credits	Credits	This field displays the credits used by this credential ID.
Reservation	Partition & Reservation	The required reservation that any job submitted has to use.
Comments	Comments & E-Mail	This field displays and allows a user to enter any comments relating to the Credential.
E-Mail Address	Comments & E-Mail	This field displays and allows a user to enter the e -mail address for the Credential.

Categories

Membership:Credential, ID, User, Group, Class, Account, Quality of Service (QoS)

Utilized Resources:Credential, ID, Utilized Job, Utilized Processors, Utilized Nodes, Utilized Processor Seconds

Soft Maximum Limits:Credential, ID, Soft Maximum Jobs,Soft Maximum Processors,Soft Maximum Nodes, Soft Maximum Processor Seconds

Hard Maximum Limits:Credential, ID, Hard Maximum Jobs,Hard Maximum Processors, Hard Maximum Nodes, Hard Maximum Processor Seconds

Priority:Credential, ID, Credential Priority

Fairshare:Credential, ID, Fairshare Type, Fairshare Target

Partition & Reservation:Credential, ID, Partition, Reservation

Statistics:Credential, ID, Enable Statistics

Credits:Credential, ID, Total Credits, Used Credits,

Default Credentials:Credential, ID, User Default, Group Default, Class Default, Account Default, Quality of Service (QoS) Default

Default Resource:Credential, ID, Partition Default, Maximum Job Default, Maximum Processor Default, Maximum Nodes Default, Maximum Processor Seconds Default

Comments & EMail:Credential, ID, Comments, E-Mail Address

4.5. Create/Modify a User Profile

Summary

Users are created by the operating system while user profiles are created by Workload Manager. When a user submits a job, that user becomes visible to Workload Manager and at that moment a credential profile is automatically created for the user.

Credential Access		
Field	**Required**	**Description**
User Name	Required	This field allows an administrator to define the name of the user. Usually this is the user's login name.
Group Access List	Not Available	The group access is defined by the operating system and cannot be defined by Workload Manager.

Credential Access		
Class Access List	Not Available	The class access is defined by the resource manager and cannot be defined by Workload Manager.
Account Access List	Optional	This field allows an administrator to define which accounts this user can access.
Default Account	Optional	This field allows an administrator to define which accounts will be automatically used if the user doesn't specify an account.
Quality of Service (QoS) Access List	Optional	This field allows an administrator to define which qualities of service (QoS) this user can access.
Default Quality of Service (QoS)	Optional	This field allows an administrator to define which quality of service (QoS) will automatically be used if the user doesn't specify a quality of service (QoS).

Resource Access		
Field	**Required**	**Description**
Partition	Optional	This field allows an administrator to define which partitions this user can access
Reservation	Optional	This field allows an administrator to define which reservation this user can access.

Fairness		
Field	**Required**	**Description**

Fairness		
Fairshare Policy	Optional	*Fairshare* is a method of enforcing cluster sharing between credentials. A *credential* is a user, group, account, class/queue, or quality of service (QoS). Fairshare tracks each credential's usage for a desired amount of time and decreases a job's start priority if the fairshare policy is violated. By decreasing a job's start priority the job will wait longer in the queue before it starts allowing other jobs to execute first. • Fairshare Floor Policy - If the user's cluster usage is below the fairshare target, then the user's start priority for the job will increase. The user's cluster usage is measured as the total percentage amount of the cluster used by the user. • Fairshare Target Policy- If the user's cluster usage is above or below the fairshare target, then the user's start priority for the job will increase or decrease accordingly. The user's cluster usage is measured as the total percentage amount of the cluster used by the user. • Fairshare Cap Policy - If the user's cluster usage is above the fairshare target, then the user's start priority for the job will decrease. The user's cluster usage is measured as the total percentage amount of the cluster used by the user. • Absolute Fairshare Policy - If a user's cluster usage exceeds the fairshare target, then the user's start priority for the job will decrease. The user's cluster usage is measured as the total number

Fairness		
Fairshare Target	Optional	This field allows an administrator to define the fairshare target for this user. Refer to the fairshare policy for an understanding of how fairshare target will be used.
Priority	Optional	This field allows an administrator to define a user's job priority. A user's job priority will increase or decrease the start priority of this user's jobs. Workload Manager, with some exceptions, will start the jobs with the highest start priority first.

Job Usage Limits		
Field	**Required**	**Description**
Maximum Executing Jobs	Optional	This field allows an administrator to set the user's maximum number of simultaneously executing jobs.
Maximum Utilized Processors	Optional	This field allows an administrator to set the user's maximum number of simultaneously utilized processors.
Maximum Utilized Processor Seconds	Optional	This field allows an administrator to set the user's maximum number of simultaneously utilized processor seconds. *Processor seconds* is defined as the number of processors utilized times the number of seconds they are utilized.
Maximum Utilized Nodes	Optional	This field allows an administrator to set the user's maximum number of simultaneously utilized nodes. A *node* is a computer consisting of 1 or more processors

General Attributes		
Field	**Required**	**Description**

General Attributes		
Comments	Optional	This field allows an administrator to enter any comments regarding the user.
Enable Statistics	Optional	This check box allows an administrator to enable or disable statistics.

Email Address		
Field	**Required**	**Description**
Email Address	Optional	This field allows an administrator to add a user's email address to Workload Manager. The email address is only for contact information and is not used by Workload Manager or the resource manager.

Credits & Charging		
Field	**Required**	**Description**
Credits	Optional	This field allows an administrator to set the total credits allocated to the user.
Used Credits	Optional (Only visible if credits have been used)	This field displays the number of credits that have been used by the user.

Usage Statistics (This is only visible if a profiles is being modified)	
Field	**Description**
Current Processor Seconds	The two charts/graphs display the number of processor seconds currently being utilized by this user compared to the total number of processor seconds currently being used by other users on the cluster. The pie chart shows the usage of this user in comparison to all the other users. The bar graph shows the average usage by this user compared to the average usage of all the other users on the cluster.

Usage Statistics (This is only visible if a profiles is being modified)	
Historical Processor Seconds	The two charts/graphs display the number of processor seconds historically utilized by this user compared to the total number of processor seconds historically used by other users on the cluster. The pie chart shows the usage of this user in comparison to all the other users. The bar graph shows the average usage by this user compared to the average usage of all the other users on the cluster.
Utilized Versus Dedicated	This line graph displays the number of processors dedicated or reserved for the user compared to the number of processors actually utilized by the user. The line graph displays the last two days of usage.

4.6. Create/Modify a Group Profile

Summary

Groups are created by the operating system while group profiles are created by Workload Manager. When a user submits a job, that user's group becomes visible to Workload Manager and at that moment, a credential profile is automatically created for the group.

Credential Access		
Field	**Required**	**Description**
Group Name	Required	This field allows an administrator to define the identification name of the group. Usually this is the login name for the group.
User Access List	Optional	This field allows an administrator to define which users can access this group.
Class Access List	Not Available	The class access is defined by the resource manager and cannot be defined by Workload Manager.
Account Access List	Optional	This field allows an administrator to define which accounts this group can access.

Credential Access		
Default Account	Optional	This field allows an administrator to define which accounts will automatically be used if the group doesn't specify an account.
Quality of Service (QoS) Access List	Optional	This field allows an administrator to define which qualities of service (QoS) this group can access.
Default Quality of Service (QoS)	Optional	This field allows an administrator to define which quality of service (QoS) will automatically be used if the group doesn't specify a quality of service (QoS).

Resource Access		
Field	**Required**	**Description**
Partition List	Optional	This field allows an administrator to define which partitions this group can access.
Reservations	Optional	This field allows an administrator to define which reservation this group can access.

Fairness		
Field	**Required**	**Description**

Fairness		
Fairshare Policy	Optional	*Fairshare* is a method of enforcing cluster sharing between credentials. A *credential* is a user, group, account, class, or quality of service (QoS). Fairshare tracks each credential's usage for a desired amount of time and decreases a job's start priority if the fairshare policy is violated. By decreasing a job's start priority, the job will wait longer in the queue before it starts, allowing other jobs to execute first. • Fairshare Floor Policy - If the group's cluster usage is below the fairshare target , then the group's start priority for the job will increase. The group's cluster usage is measured as the total percentage amount of the cluster used by the group. • Fairshare Target Policy- If the group's cluster usage is above or below the fairshare target , then the group's start priority for the job will increase or decrease accordingly. The group's cluster usage is measured as the total percentage amount of the cluster used by the group. • Fairshare Cap Policy - If the group's cluster usage is above the fairshare target , then the group's start priority for the job will decrease. The group's cluster usage is measured as the total percentage amount of the cluster used by the group. • Absolute Fairshare Policy - If a group's cluster usage exceeds the fairshare target , then the group's start priority for the job will decrease. The group's cluster usage is measured as the total number

Fairness		
Fairshare Target	Optional	This field allows an administrator to define the fairshare target for this group. Refer to the Fairshare Policy for an understanding of how fairshare target will be used.
Priority	Optional	This field allows an administrator to define a group's job priority. A group's job priority will either increase or decrease the start priority of this group's jobs. Workload Manager, with some exceptions, will start the jobs with the highest start priority first.

Job Usage Limits		
Field	**Required**	**Description**
Maximum Executing Jobs	Optional	This field allows an administrator to set the group's maximum number of simultaneously executing jobs.
Maximum Utilized Processors	Optional	This field allows an administrator to set the group's maximum number of simultaneously utilized processors.
Maximum Utilized Processor Seconds	Optional	This field allows an administrator to set the group's maximum number of simultaneously utilized processor seconds. *Processor seconds* is defined as the number of processors utilized times the number of seconds they are utilized.
Maximum Utilized Nodes	Optional	This field allows an administrator to set the group's maximum number of simultaneously utilized nodes. A *node* is a computer consisting of 1 or more processors

General Attributes		
Field	**Required**	**Description**

General Attributes		
Comments	Optional	This field allows an administrator to enter any comments regarding the group.
Enable Statistics	Optional	This check box allows an administrator to enable or disable statistics.

Credits & Charging		
Field	**Required**	**Description**
Credits	Optional	This field allows an administrator to set total credits allocated to the group.
Used Credits	Optional (Only visible if credits have been used)	This field displays the number of credits that have been used by the group.

Usage Statistics (This is only visible if a profile is being modified)	
Field	**Description**
Current Processor Seconds	The two charts/graphs display the number of processor seconds currently being utilized by this group compared to the total number of processor seconds currently being used by other groups on the cluster. The pie chart shows the relative usage of this group in comparison to all the other groups. The bar graph shows the average usage by this group compared to the average usage of all the other groups on the cluster.
Historical Processor Seconds	The two charts/graphs display the number of processor seconds historically utilized by this group compared to the total number of processor seconds historically used by other groups on the cluster. The pie chart shows the relative usage of this group in comparison to all the other groups. The bar graph shows the average usage by this group compared to the average usage of all the other groups on the cluster.
Utilized Versus Dedicated	This line graph displays the number of processors dedicated or reserved for the group compared to the number of processors actually utilized by the group. The line graph displays the last two days of usage.

4.7. Create/Modify an Account Profile

Summary

Account creation occurs in Workload Manager.

Credential Access		
Field	**Required**	**Description**
Account Name	Required	This field allows an administrator to define the identification name of the account. Usually this is the login name for the account.
User Access List	Optional	This field allows an administrator to define which users can access this account.
Group Access List	Optional	This field allows an administrator to define which groups can access this account.
Class Access List	Not Available	The Class/Queue access is defined by the resource manager and cannot be defined by Workload Manager.
Quality of Service (QoS) Access List	Optional	This field allows an administrator to define which qualities of service (QoS) this account can access.
Default Quality of Service (QoS)	Optional	This field allows an administrator to define which quality of service (QoS) will automatically be used if the account doesn't specify a quality of service (QoS).

Resource Access		
Field	**Required**	**Description**
Partition	Optional	This field allows an administrator to define which partitions this account can access.
Reservation	Optional	This field allows an administrator to define which reservation this account can access.

Fairness		
Field	**Required**	**Description**

Fairness		
Fairshare Policy	Optional	*Fairshare* is a method of enforcing cluster sharing between credentials. A *credential* is a user, group, account, class, or quality of service (QoS). Fairshare tracks each credential's usage for a desired amount of time and decreases a job's start priority if the fairshare policy is violated. By decreasing a job's start priority, the job will wait longer in the queue before it starts, allowing other jobs to execute first. • Fairshare Floor Policy - If the account's cluster usage is below the fairshare target, then the account's start priority for the job will increase. The account's cluster usage is measured as the total percentage amount of the cluster used by the account. • Fairshare Target Policy- If the account's cluster usage is above or below the fairshare target, then the account's start priority for the job will increase or decrease accordingly. The account's cluster usage is measured as the total percentage amount of the cluster used by the account. • Fairshare Cap Policy - If the account's cluster usage is above the fairshare target, then the account's start priority for the job will increase or decrease. The account's cluster usage is measured as the total percentage amount of the cluster used by the account. • Absolute Fairshare Policy - If an account's cluster usage exceeds the fairshare target, then the account's start priority for the job will be

Fairness		
Fairshare Target	Optional	This field allows an administrator to define the fairshare target for this account. Refer to the Fairshare Policy for an understanding of how fairshare target will be used.
Priority	Optional	This field allows an administrator to define an account's job priority. An account's job priority will increase or decrease the start priority of this account's jobs. Workload Manager, with some exceptions, will start the jobs with the highest start priority first.

Job Usage Limits		
Field	**Required**	**Description**
Maximum Executing Jobs	Optional	This field allows an administrator the option of setting the account's maximum number of simultaneously executing jobs.
Maximum Utilized Processors	Optional	This field allows an administrator the option of setting the account's maximum number of simultaneously utilized processors.
Maximum Utilized Processor Seconds	Optional	This field allows an administrator the option of setting the account's maximum number of simultaneously utilized processor seconds. *Processor seconds* is defined as the number of processors utilized times the number of seconds they are utilized.
Maximum Utilized Nodes	Optional	This field allows an administrator the option of setting the account's maximum number of simultaneously utilized nodes. A *node* is a computer consisting of 1 or more processors

General Attributes

General Attributes		
Field	**Required**	**Description**
Comments	Optional	This field allows an administrator the option of entering any comments regarding the account.
>Enable Statistics	Optional	This check box allows an administrator the option of enabling or disabling statistics.

Credits & Charging		
Field	**Required**	**Description**
Credits	Optional	This field allows an administrator the option of setting total credits allocated to the account.
Used Credits	Optional (Only visible if credits have been used)	This field displays the number of credits that have been used by the account.

Usage Statistics (This is only visible if a profiles is being modified)	
Field	**Description**
Current Processor Seconds	The two charts/graphs display the number of processor seconds currently being utilized by this account compared to the total number of processor seconds currently being used by other accounts on the cluster. The pie chart shows the relative usage of this account in comparison to all the other accounts. The bar graph shows the average usage by this account compared to the average usage of all the other accounts on the cluster.
Historical Processor Seconds	The two charts/graphs display the number of processor seconds historically utilized by this account compared to the total number of processor seconds historically used by other accounts on the cluster. The pie chart shows the relative usage of this account in comparison to all the other accounts. The bar graph shows the average usage by this account compared to the average usage of all the other accounts on the cluster.
Utilized Versus Dedicated	This line graph displays the number of processors dedicated or reserved for the account compared to the number of processors actually utilized by the account. The line graph displays the last two days of usage.

4.8. Create/Modify a Class Profile

Summary

Classes are created by the resource manager while class profiles are created by Workload Manager.

Credential Access		
Field	**Required**	**Description**
Class Name	Required	This field allows an administrator to define the identification name of the class. Usually this is the login name for the class.
User Access List	Optional	This field allows an administrator to define which users can access this class.
Group Access List	Not Available	This field allows an administrator to define which groups can access this class.
Account Access List	Optional	This field allows an administrator to define which accounts this class can access.
Default Account	Optional	This field allows an administrator to define which accounts will automatically be used if the class doesn't specify an account.
Quality of Service (QoS) Access List	Optional	This field allows an administrator to define which qualities of service (QoS) this class can access.
efault Quality of Service (QoS)	Optional	This field allows an administrator to define which quality of service (QoS) will automatically be used if the class doesn't specify a quality of service (QoS).

Resource Access		
Field	**Required**	**Description**
Partition	Optional	This field allows an administrator to define which partitions this class can access.
Reservation	Optional	This field allows an administrator to define which reservation this class can access.

Fairness		
Field	**Required**	**Description**

Fairness		
Fairshare Policy	Optional	*Fairshare* is a method of enforcing cluster sharing between credentials. A *credential* is a user, group, account, class, or quality of service (QoS). Fairshare tracks each credential's usage for a desired amount of time and decreases a job's start priority if the fairshare policy is violated. By decreasing a job's start priority, the job will wait longer in the queue before it starts, allowing other jobs to execute first. • Fairshare Floor Policy - If the class's cluster usage is below the fairshare target, then the class's start priority for the job will increase. The class's cluster usage is measured as the total percentage amount of the cluster used by the class. • Fairshare Target Policy- If the class's cluster usage is above or below the fairshare target, then the class's start priority for the job will increase or decrease accordingly. The class's cluster usage is measured as the total percentage amount of the cluster used by the class. • Fairshare Cap Policy - If the class's cluster usage is above the fairshare target, then the class's start priority for the job will decrease. The class's cluster usage is measured as the total percentage amount of the cluster used by the class. • Absolute Fairshare Policy - If a class's cluster usage exceeds the fairshare target, then the class's start priority for the job will decrease. The class's cluster usage is measured as the total number

Fairness		
Fairshare Target	Optional	This field allows an administrator to define the fairshare target for this class. Refer to the Fairshare Policy for an understanding of how fairshare target will be used.
Priority	Optional	This field allows an administrator to define a class's job priority. A class's job priority will increase or decrease the start priority of this class's jobs. Workload Manager, with some exceptions, will start the jobs with the highest start priority first.

Job Usage Limits		
Field	**Required**	**Description**
Maximum Executing Jobs	Optional	This field allows an administrator the option of setting the class's maximum number of simultaneously executing jobs.
Maximum Utilized Processors	Optional	This field allows an administrator the option of setting the class's maximum number of simultaneously utilized processors.
Maximum Utilized Processor Seconds	Optional	This field allows an administrator the option of setting the class's maximum number of simultaneously utilized processor seconds. *Processor seconds* is defined as the number of processors utilized times the number of seconds they are utilized.
Maximum Utilized Nodes	Optional	This field allows an administrator the option of setting the class's maximum number of simultaneously utilized nodes. A *node* is a computer consisting of 1 or more processors

General Attributes

General Attributes		
Field	**Required**	**Description**
Comments	Optional	This field allows an administrator the option of entering any comments regarding the class.
Enable Statistics	Optional	This check box allows an administrator the option of enabling or disabling statistics.

Credits & Charging		
Field	**Required**	**Description**
Credits	Optional	This field allows an administrator the option of setting total credits allocated to the class.
Used Credits	Optional (Only visible if credits have been used)	This field displays the number of credits that have been used by the class.

Usage Statistics (This is only visible if a profiles is being modified)	
Field	**Description**
Current Processor Seconds	The two charts/graphs display the number of processor seconds currently being utilized by this class compared to the total number of processor seconds currently being used by other classes on the cluster. The pie chart shows the relative usage of this classes in comparison to all the other classes. The bar graph shows the average usage by this class compared to the average usage of all the other classes on the cluster.
Historical Processor Seconds	The two charts/graphs display the number of processor seconds historically utilized by this class compared to the total number of processor seconds historically used by other classes on the cluster. The pie chart shows the relative usage of this class in comparison to all the other classes. The bar graph shows the average usage by this class compared to the average usage of all the other classes on the cluster.
Utilized Versus Dedicated	This line graph displays the number of processors dedicated or reserved for the class compared to the number of processors actually utilized by the class. The line graph displays the last two days of usage.

4.9. Create/Modify a Quality of service Profile

Summary

Quality of services are created by the operating system while quality of service profiles are created by Workload Manager. When a quality of service submits a job then that quality of service becomes visible to Workload Manager and at that moment a credential profile is automatically created for the quality of service.

Credential Access		
Field	**Required**	**Description**
Quality of service Name	Required	This field allows an administrator to define the identification name of the quality of service. Usually this is the login name for the quality of service.
User Access List	Optional	This field allows an administrator to define which users can access this quality of service.
Group Access List	Not Available	The group access is defined by the operating system and cannot be defined by Workload Manager.
Class Access List	Not Available	The Class/Queue access is defined by the resource manager and cannot be defined by Workload Manager.
Account Access List	Optional	This field allows an administrator to define which accounts this quality of service can access.
Default Account	Optional	This field allows an administrator to define which accounts will automatically be used if the quality of service doesn't specify an account.

Resource Access		
Field	**Required**	**Description**
Partition	Optional	This field allows an administrator to define which partitions this quality of service can access

Resource Access		
Required Reservation	Optional	This field allows an administrator to define which reservations jobs that access this quality of service must use.

Fairness		
Field	**Required**	**Description**

Fairness		
Fairshare Policy	Optional	Fairshare is a method of enforcing cluster sharing between credentials. A *credential* is a user, groups, account, class/queue, or quality of service (QoS). It consists of tracking each credentials usage for a desired amount of time and decreasing a jobs start priority if the *fairshare policy* is violated. By decreasing a jobs start priority the job will wait longer in the queue before it starts allowing other jobs to execute first. • Fairshare Floor Policy - If the quality of services cluster usage is below the *fairshare target* the quality of services start priority for the job will be raised. The quality of services cluster usage is measured as the total percentage amount of the cluster used by the quality of service. • Fairshare Target Policy- If the quality of services cluster usage is above or below the *fairshare target* the quality of services start priority for the job will be raised or lowered accordingly. The quality of services cluster usage is measured as the total percentage amount of the cluster used by the quality of service. • Fairshare Cap Policy - If the quality of services cluster usage is above the *fairshare target* the quality of services start priority for the job will be lowered.The quality of services cluster usage is measured as the total percentage amount of the cluster used by the quality of service. • Absolute Fairshare Policy - If a quality of services cluster

Fairness		
Fairshare Target	Optional	This field allows an administrator to define the *fairshare target* for this quality of service. Refer to the *Fairshare Policy* for an understanding of how *fairshare target*will be used.
Priority	Optional	This field allows an administrator to define a quality of services job priority. A quality of services job priority will increase or decreases the start priority of this quality of services jobs. Workload Manager, with some exceptions, will start the jobs with the highest start priority first.

Job Usage Limits		
Field	**Required**	**Description**
Maximum Executing Jobs	Optional	This field allows an administrator the option of setting the quality of services maximum number of simultaneously executing jobs.
Maximum Utilized Processors	Optional	This field allows an administrator the option of setting the quality of services maximum number of simultaneously utilized processors.
Maximum Utilized Processor Seconds	Optional	This field allows an administrator the option of setting the quality of services maximum number of simultaneously utilized processor seconds. *Processor seconds* is defined as the number of processors utilized times the number of seconds they are utilized.
Maximum Utilized Nodes	Optional	This field allows an administrator the option of setting the quality of services maximum number of simultaneously utilized nodes. A *node* is a computer consisting of 1 or more processors

General Attributes		
Field	**Required**	**Description**
Comments	Optional	This field allows an administrator the option of entering any comments regarding the quality of service.
Enable Statistics	Optional	This check box allows an administrator the option of enabling or disabling statistics.

Credits & Charging		
Field	**Required**	**Description**
Dedicated Cost	Optional	The cost to have dedicated access to this resource, regardless of whether it is being utilized.
Utilized Cost	Optional	The cost to utilize this resource.

Usage Statistics (This is only visible if a profiles is being modified)	
Field	**Description**
Current Processor Seconds	The two charts/graphs display the number of processor seconds currently being utilized by this quality of service compared to the total number of processor seconds currently being used by other quality of services on the cluster. The pie chart shows the relative usage of this quality of services in comparison to all the other quality of services. The bar graph shows the average usage by this quality of service compared to the average usage of all the other quality of services on the cluster.
Historical Processor Seconds	The two charts/graphs display the number of processor seconds historically utilized by this quality of service compared to the total number of processor seconds historically used by other quality of services on the cluster. The pie chart shows the relative usage of this quality of services in comparison to all the other quality of services. The bar graph shows the average usage by this quality of service compared to the average usage of all the other quality of services on the cluster.
Utilized Versus Dedicated	This line graph displays the number of processors dedicated or reserved for the quality of service compared to the number of processors actually utilized or used by the quality of service. The line graph displays the last two days of usage.

Quality of Service (QoS) Weight		
Field	**Required**	**Description**
Queue Time Weight	Optional	This field displays the quality of service weight factor. If a idle job is submitted to this quality of service the number of minutes that it has been in the queue will be multiplied by this value. This will the increase the jobs start priority.
Expansion Factor Weight	Optional	This field displays the quality of service weight factor. If a idle job is submitted to this quality of service it's expansion factor will be multiplied by this value. This will increase the jobs start priority.

Queue Time Threshold		
Name	**Required**	**Description**
Preemption	Optional	This field displays the number of minutes that an idle job must wait before it will be given preemptor access.
Create Reservation	Optional	This field displays the number of minutes that an idle job must wait before a job reservation will be created for it. A job reservation will guarantee it specific resources as well as a specific start time.
Resource Access	Optional	This field displays the number of minutes that an idle job must wait before it can have access to any reservations owned by the Quality of Service.

XF Threshold		
Name	**Required**	**Description**
Preemption	Optional	This field displays the expansion factor value that an idle job must be equal to or greater then before it will be given preemptor access.

XF Threshold		
Create Reservation	Optional	This field displays the expansion factor value that an idle job must be equal to or greater then before a job reservation will be created for it. A job reservation will guarantee it specific resources as well as a specific start time.
Resource Access	Optional	This field displays the expansion factor value that an idle job must be equal to or greater then before it can have access to any reservations owned by the Quality of Service.

Quality of Service Flags		
Flag Name	**Required**	**Description**
Deadline	Optional	This option will make any job with a completion deadline that is submitted to this quality of service finish before that deadline.
Dedicated	Optional	This option will make any job submitted to this quality of service to require a dedicated node. A dedicated node is a node that is completely reserved for only one job.
Enable User Reservation	Optional	This option will make any user that is a member of this quality of service able to create user/personal reservations.
Ignore All Policies	Optional	This option will make any job submitted to this quality of service exempt from all resource usage policies.
No Backfill	Optional	This option will make any job submitted to this quality of service exempt from the backfill algorithm.
No Reservation	Optional	This option will make any job submitted to this quality of service unable to create a job reservation and, therefore, only able to share resources.

Quality of Service Flags		
Next To Run	Optional	This option will make any job submitted to this quality of service run next. This is accomplished by increasing the start priority of a job to be higher then all of the other queued jobs.
Preemptee	Optional	This option will make any job submitted to this quality of service preemptable. A preemptable job can be stopped and requeued if a high priority preemptor job needs to execute.
Preemptor	Optional	This option will make any job submitted to this quality of service able to preempt any preemptable job.
Provision	Optional	This option will make any job submitted to this quality of service that requests unavailable resources such as an operating system or software to have Workload Manager setup a number of nodes with the correct resources.
Reserve Always	Optional	This option will make any job submitted to this quality of service create a reservation. Usually job reservations are created when the job starts but with this option enabled the job will create a reservation immediately.

Chapter 5. Policies

5.1. Policies Overview

Moab Workload Manager has many powerful policies that can be managed to effectively get as much workload out as possible while satisfying other desires. The policies section is intended to give control over Workload Manager's various policies.

5.2. Fairshare

Summary

Fairshare allows the cluster to be shared between different individuals and/or organizations without allowing any individual or organization the ability to monopolize the cluster. This is achieved by tracking how the cluster is used over time by each credential or user, group, class, account, and quality of service (QoS) and raising or lowering the start priorities of jobs waiting to execute. It should be noted that the start priority is used by Workload Manager to decide which jobs get executed first. The higher the start priority, the sooner a job will execute. The information collected about each credential is inserted into what is called a usage window. The length or amount of time tracked in a window is defined by the system administrator.

Often times, system administrators can not achieve the cluster sharing they desire without using multiple usage windows. Multiple usage windows allow Workload Manager to balance cluster usage differently by making the most recent window more important than more outdated windows. This is achieved by using the Usage Window Decay Factor. Essentially, the lower the decay factor, the less important outdated usage windows are.

5.2.1. Fairshare Options

- Interval Length - This field allows an administrator to define how long each window lasts.
- Depth - This field allows an administrator to define how many windows should exist.
- Usage Metric - This field allows an administrator to define how credential usage is tracked.
    1. Dedicated PS - This field tracks credential usage according to the number of processor seconds reserved for a job.
    2. Dedicated PES - This field tracks credential usage according to the number of processor equivalent seconds reserved for a job.
    3. Utilized PS - This field tracks credential usage according to the number of processor seconds used by a job.

- Decay Factor - This field allows an administrator to define how big of an influence are outdated usage windows in calculating a jobs start priority. The lower the Usage Window Decay Factor, the less important are the outdated usage windows.

5.2.2. Fairshare Table

- Credential Type - This field displays the credential type whether it be a user, group, account, class, or quality of service (QoS).
- Credential ID - This field displays the credential's ID.
- Fairshare Credentials Policy - This field displays the fairshare credentials policy. Consult the create user, group, class, account, or quality of service (QoS) documentation for more information regarding the policies.
- Fairshare Credentials Target - This field displays the fairshare credentials policy. Consult the create user, group, class, account, or quality of service (QoS) documentation for more information regarding the targets.
- Percentage Cluster Usage - This field displays the percentage of the cluster that was used by this credential ID in comparison to the other credential IDs for this Credential Type.
- Current Interval - This field displays the first usage window. The decay factor does not affect this window at all.
- Interval 1 - 31 - This field displays the usage interval windows 1 through 31. The decay factor affects these windows with the most outdated window being window 31 and the most recent window being window 1.

5.3. Priority

Summary:

A job has one start priority which is used to decide when a job will start. The higher the start priority the sooner a job will start. The job start priority can be anywhere between 1,000,000,000 and -1,000,000,000.

How to read priorities.

A job has one start priority which is used to decide when a job will start. The higher the start priority the sooner a job will begin execution. Workload Manager uses the priority policies to calculate a job's start priority.

A subcomponent priority of 0 means the subcomponent will be ignored. A positive subcomponent priority means the start priority will be increased. A negative subcomponent priority means the start priority will be decreased. Refer to the documentation below for information about Main Components.

How to understand the priority window layout.

Workload Manager uses 39 components to calculate the start priority. These components are grouped into tabs according to their functionality. The Main Component is different from the subcomponents. Refer to the documentation below for further information about Main Components.

The table shows only idle/queued jobs and their start priority. Only idle/queued jobs are displayed because priority policies do not affect running jobs.

What are components and subcomponents?

The 7 component groupings are crucial to understanding priorities. The Main Component tab is used only to increase or decrease the subcomponents priorities. The subcomponents increase or decrease the job start priority.

How the job start priority is calculated.

A job has one start priority which is used to decide when a job will start. The higher the start priority the sooner a job will start. Workload Manager uses the priority policies to calculate a job's start priority.

The start priority is calculated first by adding all the subcomponents in a group together and multiplying the total of these subcomponents by the Main Component priority. This process is repeated 7 times. Once for each main component. The start priority is a summation of these 7 totals.

It should be noted that if the Main Component priority is set to 0, all of the subcomponent priorities for that Main Component will be ignored.

How to enable a priority

To enable a priority, two priorities must changed. The first priority is the sub component priority and the second priority is the Main Component priority.

For example to apply a priority of 1 for a user's priority, the user priority in the subcomponent credential's priorities must be set to 1 and the Credential Priorities in the Main Component must also be set to 1.

What does a 0 mean?

If the Main Component priority is set to 0, all of the subcomponent priorities for that Main Component will be ignored. A sub component priority of 0 means the sub component will be ignored.

5.3.1. Main Priority Components

- Wait Time Job Services - This field allows an administrator to increase or decrease all of the Wait Time Job Services priorities. If this is set to 0 all of the subcomponents prioritie's for Wait Time Job Services will be ignored.
- Wait Time Job Targets - This field allows an administrator to increase or decrease all of the Wait Time Job Target's priorities. If this is set to 0 all of the subcomponents priorities for Wait Time Job Targets will be ignored.
- Fairshare Usage - This field allows an administrator to increase or decrease all of the Fairshare Usage priorities. If this is set to 0 all of the subcomponents priorities for Fairshare Usage will be ignored.
- Resource Requests - This field allows an administrator to increase or decrease all of the Resource Request's priorities. If this is set to 0 all of the subcomponents priorities for Resource Request's will be ignored.
- Credential Priorities - This field allows an administrator to increase or decrease all of the Credential Priority's priorities. If this is set to 0 all of the subcomponents priorities for Credential Prioritie's will be ignored.
- Job Attributes - This field allows an administrator to increase or decrease all of the Job Attributes priorities. If this is set to 0 all of the subcomponents priorities for Job Attributes will be ignored.
- Executing Job Usage - This field allows an administrator to increase or decrease all of the Executing Job Usage priorities. If this is set to 0 all of the subcomponents priorities for Executing Job Usage will be ignored. Unlike the other components, this component only effects executing jobs and is only applicable when preemption is used.

5.3.2. Wait Time Job Services

- Queue Time - This field allows an administrator to set the priority of a job according to the minutes the job has waited in the queue.
- Expansion Factor (X-Factor) - This field allows an administrator to set the priority of a job according to the expansion factor* of the job.
- Policy Violation - This field allows an administrator to set the priority of a job according to whether the job has violated a usage limit. If the job has violated a usage limit, the job is assigned a policy violation value of 1; otherwise the job is assigned a policy violation value of 0.
- By Pass - This field allows an administrator to set the priority of a job according to the number of other jobs that have started execution before this job. The other jobs are only counted if Workload Manager started the other jobs because of a backfill policy.
- Dead Line - This field allows an administrator to set the priority of a job according to the proximity of the job's deadline. The closer to the proximity, the higher the dead line value.

5.3.3. QoS Targets

- Queue Time - This field allows an administrator to set the priority of a job according to Quality of Service queue time target. The closer the job is to this target, the higher the Queue Time value.

- Expansion Factor (X-Factor) - This field allows an administrator to set the priority of a job according to Quality of Service expansion factor target. The closer the job is to this target, the higher the Expansion Factor value.

5.3.4. Credential Priority

- User - This field allows an administrator to set the priority of a job according the User's priority.
- Group - This field allows an administrator to set the priority of a job according the Group's priority.
- Account - This field allows an administrator to set the priority of a job according the Account's priority.
- Class - This field allows an administrator to set the priority of a job according the Class priority.
- QoS - This field allows an administrator to set the priority of a job according the Quality of Service (QoS) priority.

5.3.5. Job State Priority

- Job Attribute - This field allows an administrator to set the priority of a job according a job's attributes. Refer to the Workload Manager Priority Factors documentation for information on how to set the Job Attributes.
- Job State - This field allows an administrator to set the priority of a job according a job's state. Refer to the Workload Manager Priority Factors documentation for information on how to set the Job state.

5.3.6. Fairshare Usage Priority

- User - This field allows an administrator to set the users priority of a job according to the fairshare usage.
- Group - This field allows an administrator to set the groups priority of a job according to the fairshare usage.
- Account - This field allows an administrator to set the accounts priority of a job according to the fairshare usage.
- Class - This field allows an administrator to set the classes priority of a job according to the fairshare usage.
- Quality of Service (QoS) - This field allows an administrator to set the qualities of service (QoS) priority of a job according to the fairshare usage.
- Jobs Per User - This field allows an administrator to set the priority of a job according to the number of jobs currently executing for this user.
- Processor Seconds - This field allows an administrator to set the priority of a job according to the number of processor seconds currently being used by this user.

- Processors Per User - This field allows an administrator to set the priority of a job according to the number of processors currently being used by this user.

5.3.7. Resource Requests Priority

- Node - This field allows an administrator to set the priority of a job according to the total number of nodes requested by the job. The more nodes requested, the higher the Node value.
- Disk - This field allows an administrator to set the priority of a job according to the total amount of disk space requested by the job. The more disk space requested, the higher the Disk value.
- Processor - This field allows an administrator to set the priority of a job according to the total number of processors requested by the job. The more processors requested, the higher the Processor value.
- Memory - This field allows an administrator to set the priority of a job according to the total amount of memory in megabytes requested by the job. The more memory requested, the higher the Memory value.
- Swap - This field allows an administrator to set the priority of a job according to the total amount of swap in megabytes requested by the job. The more swap requested, the higher the Swap value.
- Processor-Seconds - This field allows an administrator to set the priority of a job according to the total number of processor seconds requested by the job. The more processor seconds requested, the higher the Processor Seconds value.
- Processor-Equivalent - This field allows an administrator to set the priority of a job according to the total number of processor equivalents requested by the job. The more processor equivalents requested, the higher the Processor Equivalents value.
- Wall Time - This field allows an administrator to set the priority of a job according to the total amount of wall time seconds requested by the job. The more wall time requested, the higher the Wall Time value.

5.3.8. Executing Job Usage Priority

- Consumed - This field allows an administrator to set the priority of a job according to the total number of processor seconds it has consumed. Unlike other components, this component only effects executing jobs and is only applicable when preemption is used.
- Hunger - This field allows an administrator to set the priority of a job according to the total number of processors needed to balance a dynamic job's performance. This field is only applied to dynamic jobs.
- Remaining - This field allows an administrator to set the priority of a job according to the total number of processor seconds it has remaining. Unlike other components this component only effects executing jobs and is only applicable when preemption is used.
- Percentage Consumed - This field allows an administrator to set the priority of a job according to the percentage of the wall clock that has been consumed. Unlike other components this component only effects executing jobs and is only applicable when preemption is used.

5.3.9. Priority Display Options

- View Subcomponents in Table as actual values. - This option displays the actual subcomponent values for the jobs in the table.
- View Subcomponents in Table as percentage values. - This option displays the subcomponent percentage breakdown for the subcomponent's group.
- Display start priority pie chart - This option displays a pie chart of the priority components. Negative components are not displayed.
- Display start priority bar graph - This option displays a bar chart of the priority components. Positive and negative components are displayed.

5.4. Job Policies and Settings

Contained in this section:

5.4.1. Job Limit Settings

Summary

This window is used to place system wide restrictions on jobs.

- Wall Clock - This field allows an administrator to specify the maximum amount of wall clock time that can be requested by any single job.
- Maximum Job Start - This field allows an administrator to specify the maximum number of times Workload Manager will attempt to start the job.
- Maximum Job Preempt - This field allows an administrator to define the maximum number of times a job can be preempted by Workload Manager for higher priority jobs.
- Maximum Processors - This field allows an administrator to define the maximum number of processors that can be requested by any jobs.
- Maximum Processor Seconds - This field allows an administrator to define the maximum number of processor seconds that can be requested by any job. Processor seconds are defined as the number of processors used by a job times how long the job executed.
- Exceeded Wallclock Job Violation - This field allows an administrator to define the amount of time Workload Manager will allow a job to exceed its wallclock limit before it is terminated.

5.4.2. Job Defer Settings

Summary

- Defer Wait Time - This field allows an administrator to define the amount of time a job will be held in the deferred state before being released back to the idle job queue.
- Starts Before Defer - This field allows an administrator to define the amount of time a job will be allowed to fail in its start attempts before being deferred.
- Defers Before Hold - This field allows an administrator to define the the number of times a job can be deferred before it will be placed in batch hold.
- Synchronization Wait - This field allows an administrator to define the length of time after which Workload Manager will change a job's expected state to an unexpected reported state. It should be noted that Workload Manager will not allow a job to run as long as its expected state does not match the state reported by the resource manager.

5.4.3. Global Job Policy Settings

Summary

This window contains job-specific global settings.

- Job Priority Policy - This field allows an administrator to specify when a job's start priority should increase. With some exceptions, the higher a job's start priority, the sooner the job will start.
    1. Always - This policy will begin increasing a job's start priority relative to the time it has waited to execute.
    2. Full Policy - This policy will begin increasing a job's start priority once all the usage violations have vanished.
    3. Queue Policy - This policy will begin increasing a job's start priority once all the queue usage violations have vanished.
- Use Machine Speed - By checking this box, an administrator specifies that a job's wall clock should be increased if the job is executing on a slower node and that a job's wall clock should be decreased if the job is executing on a faster node. The speed of the node is assessed by examining the node speed option located in the list nodes window.

5.5. Reservation Policies

Summary

Workload Manager uses reservations to guarantee that a specific amount of resources will be available for a given job or set of users at a particular time. For example, Workload Manager can reserve 20 processors and 10 GB of memory for users Bob and John from Friday 6:00 AM to Saturday 10:00 PM. Workload Manager uses reservations internally to manage backfill, protect job resources, allow service

guarantees, support deadlines and QoS, and enable grid scheduling. Workload Manager supports infinite, reoccurring, and one-time reservations.

When backfill is enabled, Workload Manager will attempt to schedule lower-priority jobs ahead of a higher-priority job that can't start immediately. In order to ensure that those low-priority jobs don't delay the high-priority job's start time, Workload Manager can reserve the resources needed by the high-priority job. These are called *priority reservations*. The reservation policy determines how Workload Manager handles priority reservations.

Reservation Settings		
Field	***Required***	***Field Information***
Policy	Required	This is the policy Workload Manager uses when creating priority reservations. These reservations protect the resources a job is using until the job completes. • *CurrentHighest* - Existent priority reservations will be relinquished to new jobs with higher priority • *Highest* - All idle jobs that receive a reservation will keep it until they run even if new jobs are higher priority • *Never* - No idle jobs receive reservations
Depth	Required	This is how many priority reservations Workload Manager will create. A higher value will protect the start time of high-priority jobs, but may decrease backfill efficiency.
Reservations Per Node	Required	This is the maximum number of priority reservations that can be created on any single node. On large SMP systems, this value should be set to approximately twice the number of reservations that exist on the system.

Reservation Settings		
Retry Time	Optional	This is the period of time Workload Manager attempts to re-start a job that received an priority reservation that originally failed to start.
Creation Policy	Optional	This determines which users can create one-time reservations (also called *Administrative* reservations). This setting is unrelated to priority reservations.

5.6. Resource Violation

Summary

The resource violation policies dictate how Workload Manager will handle jobs that use more resources than they request. Workload Manager monitors a job's usage of processors, disk space, swap space, and memory. If a job exceeds its allocation for one of these resources, Workload Manager can be configured to take one of several actions under several different violation policies.

Resource Violation Settings		
Field	***Required***	***Field Information***

Resource Violation Settings		
Policy	Required	This determines what action Workload Manager will take when it detects a resource violation. • Never - No action is taken • Always - An action is taken immediately upon detecting a violation • ExtendedViolation - An action is taken only if a detected violation persists for more than the specified time limit • BlockedWorkloadOnly - Considers all possible combinations of jobs that can run on the available resources and selects the best combination (see the *Attribute* parameter below)
Action	Optional	This is the number of jobs in the queue Workload Manager should consider for backfill. By default, all jobs are considered. If *Depth* is set, Workload Manager will only consider that number of jobs for backfill scheduling. For example, if there are 15 idle jobs in the queue and *Depth* is set to 10, only 10 jobs would be considered for backfill. If there are fewer than 10 jobs in the queue, all will be considered. Setting this number higher will result in a higher utilization and better turn-around times, especially for smaller jobs; but may result in low-priority jobs being started before medium-priority jobs. This parameter should be tuned for your specific situation.

Resource Violation Settings		
Attribute	Optional	This is the criteria used by the backfill algorithm to determine the best jobs to backfill. For example, if Procs is selected, a job that requires the exact amount of available processors will be considered the best. This parameter only applies to the BestFit and Greedy backfill policies. • *Procs* - This is the number of processors. • *ProcSeconds* - This is the number of processors multiplied by the duration of the job, in seconds. • *Seconds* - This is the duration (or wallclock time) of the job, in seconds. • *PE* - This is the processor equivalence of a job (see explanation below). • *PESeconds* - This is the processor equivalence of a job multiplied by the duration of the job, in seconds.

5.7. Node Policies

Summary

Node Task Allocation Policy

A *task* is a request for resources that must exist on a single compute node. Each job may have one or more tasks. Workload Manager allocates resources to jobs based on the tasks in the job. This is useful because nodes with multiple processors are usually able to support more than one task at a time. For example, if a job has 2 tasks where each task requires 1 processor and 256 MB of memory, Workload

Manager may choose to allocate the job to a dual processor node with 512 MB of memory, or to 2 single processor nodes with 256 MB of memory each. The node-task allocation policy determines which tasks may run on the same node.

Node Task Allocation		
Field	***Required***	***Field Information***
Policy	Required	This is the algorithm Workload Manager uses to determine which tasks may run on the same node. • Shared - Tasks from any job and any user may use available resources on any node. • SingleUser - For any given node, only tasks from jobs submitted by the same user may run. • SingleJob - Only tasks from the same job may run for any given node. • SingleTask - Only one task may run on each node.

Node Availability Policy

Workload Manager will start jobs on nodes that are not full and are not considered busy. Workload Manager considers a node busy according to which Node Availability Policy is set .

Node Availability Policy		
Field	***Required***	***Field Information***

Node Availability Policy		
Policy	Required	This is the algorithm Workload Manager uses to determine if a node is busy. • Utilized - The utilized (in use) resources on the node equal the configured (total available) resources. • Dedicated - The dedicated (assigned or reserved) resources on the node equal or exceed the configured resources. • Combined - Either of the above two conditions is met. Resources may be dedicated to a user, group, or account for a specific period of time. But some of those resources may not be used during the entire period. This setting allows Workload Manager to differentiate between the two possibilities.

Node Allocation Policy

Node allocation is the process of selecting the best resources, from a list of available resources, to assign to a job. Making this decision intelligently is important in environments with heterogeneous resources or nodes that can support multiple jobs at the same time.

Node Allocation Policy		
Field	***Required***	***Field Information***

Node Allocation Policy		
Policy	Required	This is the algorithm Workload Manager uses to allocate nodes. • CPULoad - Nodes that have the maximum amount of available, unused CPU power are selected. This is good for timesharing systems, but is only applicable to jobs starting immediately. For future jobs, the MinResource policy is used. • FirstAvailable - Nodes are allocated in the order they are presented by the resource manager. • LastAvailable - Resources are selected so as to minimize the amount of time the resources remain unused after the job completes. This minimizes node-time fragmentation and is useful in systems that have a large number of reservations. • MinResource - Nodes that have the smallest amount of resources that meet the job's requirements are selected. • Contiguous - Nodes are allocated in contiguous (linear) blocks. This is required by the Compaq RMS system. • MaxBalance - Nodes that are as similar as possible to each other are allocated to each job. The most important consideration in determining node similarity is node speed. • Fastest - The fastest available nodes are allocated to each job. Workload Manager determines which nodes are fastest based upon first the node speed and then the processor speed of each node. If neither of these values is available, the nodes

Node Allocation Policy

5.8. Partition Policies

Summary:

This section deals with policies relating partitions and their behavior. Below is a list of partition policies.

5.8.1. Partition Allocation Policy

A direct way to assign a peer allocation algorithm when multiple partitions are available for a job. Because clusters are considered partitions, this defines how jobs can be migrated to remote resources if multiple remote clusters can be found. Values and their descriptions are listed in the table below.

- BESTFIT - Allocate resources from the eligible peer with the fewest available resources - measured in tasks (minimizes fragmentation of large resource blocks).
- BESTFITP - Allocate resources from the eligible peer with the fewest available resources - measured in percent of configured resources (minimizes fragmentation of large resource blocks).
- FIRSTFIT - Allocate resources from the eligible peer which can start the job the soonest.
- FIRSTCOMPLETION - Allocate resources from the eligible peer which can complete the job the soonest (takes into account data staging time and job-specific machine speed.
- LOADBALANCE - Allocate resources from the eligible peer with the most available resources - measured in tasks (balances workload distribution across potential peers).
- LOADBALANCEP - Allocate resources from the eligible peer with the most available resources - measured in percent of configured resources (balances workload distribution across potential peers).
- ROUNDROBIN - Allocate resources from the eligible peer which has been least recently allocated.

5.9. Backfill

Summary

Backfill is an optimization policy that allows a scheduler to make better use of available resources by running jobs out of order. When using Backfill, Workload Manager prioritizes the jobs in the queue into a sorted list with the highest priority job first. Beginning at the top of the list, it starts the jobs one by one until it reaches a job that it cannot start because the necessary resources are not available. Using the start times and wall clock limits of the currently running jobs, Workload Manager then calculates when it will be able to start the job. It reserves that spot (in the future) for the job, and attempts to schedule some of

the remaining lower-priority jobs in the gaps left over from the higher-priority jobs. This process continues until Workload Manager has attempted to start all the jobs in the list, until all resources are consumed, or until Workload Manager has considered a specific number of jobs. Backfill allows Workload Manager to achieve a higher utilization than would be otherwise possible, while remaining mostly true to the original job priorities.

Backfill Settings		
Field	***Required***	***Field Information***
Policy	Required	In this field you specify the kind of backfill algorithm Workload Manager uses to schedule jobs. • *None* - Backfill is not enabled • *FirstFit* - Considers jobs in the queue sequentially, beginning with the highest priority and moving down the list • *BestFit* - Considers all jobs in the queue and selects the job that best fits the available resources (see the *Attribute* parameter below) • *Greedy* - Considers all possible combinations of jobs that can run on the available resources and selects the best combination (see the *Attribute*parameter below)

Backfill Settings		
Depth	Optional	This is the number of jobs in the queue Workload Manager should consider for backfill. By default, all jobs are considered. If *Depth* is set, Workload Manager will only consider that number of jobs for backfill scheduling. For example, if there are idle jobs in the queue and *Depth* is set to 10, only 10 jobs would be considered for backfill. If there are fewer than 10 jobs in the queue, all will be considered. Setting this number higher will result in a higher utilization and better turn-around times, especially for smaller jobs, but may result in low-priority jobs being started before medium-priority jobs. This parameter should be tuned for your specific situation.

Backfill Settings		
Attribute	Optional	This is the criteria used by the backfill algorithm to determine the best jobs to backfill. For example, if *Procs* is selected, a job that requires the exact amount of available processors will be considered the best. This parameter only applies to the *BestFit* and *Greedy* backfill policies. • *Procs* - This is the number of processors • *ProcSeconds* - This is the number of processors multiplied by the duration of the job, in seconds • *Seconds* - This is the duration (or wallclock time) of the job, in seconds • *PE* - This is the processor equivalence of a job (see explanation below) • *PESeconds* - This is the processor equivalence of a job multiplied by the duration of the job, in seconds

5.10. Role Based Authorization

Summary

The role based authorization policies dictate what commands may be run by each level of administrator. These settings are saved in Workload Manager and affect command-line operations as well as permissions within Cluster Manager.

Role Based Authorization Settings	
Field	***Field Information***

Role Based Authorization Settings	
Name	name or short description of the role
showstate	show current state of resources
showq	show queued jobs
setspri	adjust job priority or system priority of job (Maui compatibility)
setres	set an admin or user reservation (Maui compatibility)
sethold	set job holds (Maui compatibility)
releasehold	release job defers and holds (Maui compatibility)
showstats	show scheduler usage statistics
resetstats	reset scheduler usage statistics
releaseres	release reservations (Maui compatibility)
showres	show existing reservations
diagnose	provide diagnostic report for various aspects of resources, workload, and scheduling (Maui compatibility)
showstart	show estimates of when job can or will start
setqos	modify job QOS settings (Maui compatibility)
showbf	show current resource availability
showconfig	show current scheduler configuration (Maui compatibility)
checkjob	provide detailed status report for specified job
checknode	provide detailed status report for specified node
runjob	force a job to run immediately (Maui compatibility)
canceljob	cancel job (Maui compatibility)
changeparam	change in memory parameter settings (Maui compatibility)
mjobctl	control and modify jobs
mnodectl	control and modify nodes
mrsvctl	control and modify reservations
mschedctl	modify scheduler state and behavior
mdiag	provide diagnostic report for various aspects of resources, workload, and scheduling
mshow	display various diagnostic messages about the system resources and job queues
mbal	execute a command on a remote machine chosen by a load-balancing algorithm
mcredctl	control and modify scheduler credential objects
mrmctl	control and modify resource managers

Role Based Authorization Settings	
msub	submit a job directly for migration to an appropriate resource manager

Chapter 6. Statistics

6.1. Statistics Overview

Cluster Manager offers a wide assortment of customizable statistics whether they be quick charts, customized charts, graphs, and reports, or estimation matrix statistics.

6.2. Quick Charts/Graphs

Quick Charts provides a simple interface for viewing the most common statistics. Statistics are gathered from the first day of the current month to the last day of the current day, week, or month.

Available Charts

1. System Overview
2. Total Processor Hours Per Account
3. Queue Time Per Account
4. Total Processor Hours Per User
5. Resource Requests Per User
6. Queue Time Per Quality of Service (QoS)

6.3. Matrix Statistics

Summary

Matrix statistics are used both to analyze historic workload and predict future workload. The left column of the table displays different job processor sizes. The top column displays relative time frames in the format Hours:Minutes:Seconds.

For example, the table below would be understood as follows. For jobs using 1 processor, 12 jobs completed within the first 15 minutes of execution and 8 jobs completed after the first 15 minutes and before the first hour of execution. For jobs using 4 processors, 5 jobs completed within the first 15 minutes of execution and 0 jobs completed after the first 15 minutes and before the first hour of execution.

Total Completed Jobs	00:15:00	01:00:00
1 Processor	12	8

4 Processors	5	-

Matrix Statistics Types

- Estimated Start Time - This field displays the predicted start time of a created/submitted job according to the number of processors the job would use. This information can help users determine how many processors they should submit a job to for optimal start time. For example, it may take less time to start a four-hour job submitted to four processors, than to one processor for a sixteen-hour job.
- Average Expansion Factor - This field displays the historic average expansion factor of a job according to the number of processors it used. The expansion factor is calculated using the following equation: ((queue time of a job + job's duration) / job duration).
- Maximum Expansion Factor - This field displays the historic maximum job expansion factor of a job according to the number of processors it used. The expansion factor is calculated using the following equation: ((queue time of a job + job's duration) / job duration).
- Average Queue Time - This field displays the historic average wait time before a job starts executing, according to the number of processors it used. Queue time is the number of hours a job waited before it began execution.
- Average Bypass - This field displays the historic average by pass of a job according to the number of processors it used. By pass is the number of jobs that started execution before this job because of backfill policies. This is useful in recognizing which type of jobs are being by passed by backfill.
- Maximum Bypass - This field displays the historic maximum by pass of a job according to the number of processors it used. By pass is the number of jobs that started execution before this job because of backfill policies. This is useful in recognizing which type of jobs are being by passed by backfill.
- Total Completed Jobs - This field displays the total number of jobs that completed in the time interval according to the number of processors they used.
- Cluster Processor Hours Requested (%) - This field displays a breakdown of the requested time on the cluster according to the number of processor hours. Processor hours are the number of processors times the number of hours that they were requested. The sum total of the table is 100%. Each cell inside the table gives the percentage of the total cluster processor hours requested by jobs of that size and duration.
- Cluster Processor Hours Utilized (%) - This field displays a breakdown of the utilized time on the cluster according to the number of processor hours. Processor hours are the number of processors times the number of hours that they were utilized. The sum total of the entire table is 100%. Each cell inside the table gives the percentage of the total cluster processor hours utilized by jobs of that size and duration.
- Wall Clock Accuracy (%) - This field displays the average wall clock accuracy, or user estimate accuracy of how long a job would execute, according to the number of processors it used. A value greater than 100 indicates the average user overestimates the job wall clock time. A value less then 100 indicates the average user underestimates the job wall clock time. A value of 100 indicates the average user estimates the job wall clock accurately.
- Backfill Count (%) - This field displays the percentage of jobs that were delayed in executing because the backfill policy made them execute later.

- Backfill Processor Hours Utilized (%) - This field displays the percentage of processor hours for jobs that were delayed in executing because of the backfill policy that later executed.
- Job Efficiency (%) - This field displays the average percentage of the CPU that jobs used according to the number of processors of each job.
- Quality of Service (QoS) Delivered (%) - This field displays the average percentage of jobs that received their desired quality of service (QoS) according to the number of processors they used.

6.4. Custom Charts/Graphs

Summary

This window allows one to create charts and graphs showing statistics over a custom time interval.

6.4.1. Credential Based Charts

Select Credentials

The drop down box allows you to select which credential the statistics will be calculated for or if the cluster-wide statistics should be used. The Display All Credentials option will display all the credentials that have been tracked by the statistics regardless of whether they have any activity recorded. The Display Credentials With Statistics option will display only the credentials that have recorded some type of activity. The Display Listed Credentials option displays only the requested credentials.

Select Criteria

- Execute Jobs - This field displays only completed jobs or jobs that have finished execution. If the line graph is selected, the resulting points are the number of jobs that completed at that exact moment of time.
- Total Processor Hours - This field displays the number of hours used on the cluster. Processor hours are defined as the number of processors used times how long each was used. For example, a user who uses 5 processors for 5 hours would have used 25 processor hours. However,a user who used 1 processor for 5 hours would only have used 5 processor hours.
- System Utilization* - This field displays the number of processors used by the job.
- Total Queue Time** - This field displays the total hours a credential's jobs waited in the queue before starting.
- Total Backlog** - This field displays the backlog. The backlog is the number of processor hours a job waited in the queue. Processor hours are defined as the number of processors used times how long each was used.
- Total Expansion Factor** - This field displays the expansion factor. Expansion factor is defined as ((Queue Time + Execution Time) / Wall Clock).

- Resource Requests - This field displays three fields. Processor-seconds, memory, and wall clock per credentials that either exceeded the requested resource or under-utilized the requested resource.
- Jobs That Met QOS Target** - This field displays the percentage of jobs that met their QOS target to total jobs.
- Allocated Nodes*** - This field displays the number of nodes allocated to this credential
- Allocated Processors*** - This field displays the number of processors allocated to this credential

Note: *Note: Because Workload Manager operates by averaging usage across it's statistical intervals the values displayed can be misleading. For example, if the statistical interval for Workload Manager was set to 10 minutes for a cluster of 256 processors and one job, which used all 256 processors, was submitted that started and ended in 5 minutes, then an administrator would assume the System Utilization would display 100% of the processors used. However, because the interval was 10 minutes long and the job only ran for 5 minutes, the average System Utilization for the 10 minute interval was 50%.

Note: **Note: This field can only be calculated when a job finishes execution.

Note: ***Note: Requires that Moab is running under a dedicated node model, where a node can be running only 1 job at at time.

Chart Title

The text in the chart title field will be displayed on the top of the chart/graph. By default, the title is the selected criteria, but the title can be edited to match personal preferences.

Display Flags

The pie charts and bar graphs can be customized to display data in different manners.

1. If the Average check box is selected, the charts will display the average value over the requested time frame.
2. If the Maximum check box is selected, the charts will display the maximum value that occurred over the requested time frame.
3. If the Total check box is selected, the charts will display the total value over the requested time frame.

It should be noted that line graphs and the Resource Request field do not support display flags.

Chart Type

For formatting, you can choose from Pie Chart, 3D Pie Chart, Bar Graph, 3D Bar Graph, and Line Graph.

Time Frame

Choose a time frame for the graph. Time frames can be chosen from the basis of Month, Week, Day, Hour or Custom. The Month time frame gathers data from the first of the month to the end of the month. The Week time frame gathers data from the start of the week to the end of the week. The Day time frame gathers data from the start of the day to the end of the day. The Hour time frame gathers data started from the selected hour and ending one hour from that time. The Custom time frame gathers data from the start time and ends at the end time.

6.4.2. Node Categorization Charts

Summary

Create charts that show node categorization over time. For example, one might create a line chart that will show when nodes were in a hardware failure state or create a bar graph to show how much time nodes spent in user reservations.

Creating a stacked line graph brings up a chart window that allows one to dynamically make node categories visible or invisible. It shows these node categories in a tree structure where branches corresponding to node categories can be expanded or contracted. As these node categories are expanded or contracted the chart is immediately updated. For example, one might collapse the hardware failure and software failure categories into the parent down time node category.

Why do only 4 of the node categories appear on the chart?

By default, Moab will only categorize nodes into the states idle, active, hardwareFailure, and [NONE]/Other. Usually [NONE]/Other means that Moab was not running. In order to see the other node categories, one must set a reservation on the node indicating the desired category. For example, for node001 to appear as being in the hardware maintenance state one would set a reservation on node001 whose duration matched the length of the hardware maintenance. Administrators wanting accurate node categorization charts should consider using triggers to set reservations when nodes change state.

6.4.3. Generic Metric Charts

Summary

This window allows one to create chart that show generic metrics over time

Details

To create a chart, simply select the desired nodes and generic metric and click the create button. A chart should appear. If there are too many nodes in your system to view in a single chart, the paging buttons will be enabled. Click the "Next Page" button to see the next 10 nodes.

One can also filter out data sequences which are not of interest. To do this, click the filters checkbox at the lower left of the chart. A lower panel should appear allowing you to select nodes that either have a value or average above or below a certain threshold. Clicking the "Apply Filter" button causes this change to be reflected in the chart.

What to do if you see a warning that generic metrics aren't configured

If a chart cannot be created due to a warning informing you that no generic metrics are configured in Moab Workload Manager, you should check to see that your resource manager is returning generic metric information to Moab Workload Manager. Generic Metrics are usually returned to Moab Workload Manager through the CLUSTERQUERYURL configured in your moab.cfg.

6.4.4. Job Template Charts

Summary

This window allows one to create chart that show statistics relating to job templates

Details

To create a chart, simply select the desired job templates, usage metric, calculation parameters, chart type, and time frame and click the create button.

Why am I warned that no job templates are configured?

If you cannot create a chart due to a warning you that statistics for job templates is not configured in Moab Workload Manager you need to modify your moab.cfg. Most likely you either have no job templates configured or you did not add the JSTAT attribute. Moab Workload Manager will not collect statistical data for job templates not configured with a JSTAT attribute. For more information on configuring job templates with JSTAT, see the Moab Administrators Guide.

6.5. Custom Reports

This window is used to generate reports about credentials on the cluster. It should be noted that statistics will not be recorded, and, consequently, statistics will not be available for reports, if statistic tracking for credentials have not been enabled in Moab Workload Manager.

Basic Report

This report will display only the selected credentials and their usage according to the specified usage metric.

Advanced Report

This report will display the children credentials and their usage according to the specified usage metric. The children credentials will be reordered and displayed according to which parent credential they are associated with.

Detailed Summary Report

This report will multiple calculations about a credential's activity over the specified time frame

Time Frame

Choose a time frame for the graph. Time frames can be chosen from the basis of wonth, week, day, hour, or custom. The Month time frame gathers data from the first of the month to the end of the month. The Week time frame gathers data from the start of the week to the end of the week. The Day time frame gathers data from the start of the day to the end of the day. The Hour time frame gathers data started from the selected hour and ending one hour from that time. The Custom time frame gathers data from the start time and ends at the end time.

Chapter 7. Diagnostics

7.1. Diagnostics Overview

Diagnostics are intended to give an adminstrator quick and easy system information for diagnosing potential problems.

7.2. Diagnostics Support

Summary:

The diagnostics support features allows the user to run a set of commands that will check the status of various parts of their system. These commands are controlled by a script named, "support.diag.pl". This should be in your $WORKLOADHOMEDIR/tools directory or wherever your moab tools directory has been installed.

> **Warning**
>
> This window will not work without this script.

The diagnostics support screen is intended to allow the user to select from the tree what commands to run using a built in script. The script will then package the output of each command into a file. This file is saved in an output directory specified by the user or by default in "/tmp/".

If anything should go wrong in the support process, a pop box will open saying what happened. If this is not sufficient, please consult the mcm.log file for more information.

Chapter 8. Miscellaneous

8.1. Miscellaneous Overview

Various Cluster Manager sections that don't fit in other categories are contained here.

8.2. Console

Cluster Manager communicates directly with Workload Manager. This console displays the commands submitted to Workload Manager from Cluster Manager, as well as any information returned by Workload Manager. Workload Manager output messages will be highlighted in green while error messages will be highlighted in red.

- Automatically Process Commands - This field, when enabled, will automatically submit each command to Workload Manager and place the command and results in the Output text window. If this field is not enabled, the commands that were to be submitted to Workload Manager will be placed in the Commands text window.
- Process Commands - This field will submit to Workload Manager any text in the Commands text window.

8.3. Debugging and Log Levels

Allows users to select the log level in Moab Cluster Manager, which can be used to help prepare logs to be accompanied with bug reports. Logs will be written to the <MCM_HOME>/logs/mcm.log file. Higher logging levels create more detailed logging information, which facilitates debugging but may slow performance.

Below are the logging levels available, listed in order of increasing verbosity.

- 0 - Off : Turns off logging.
- 1 - Fatal : Logs only server events that cause the application to abort.
- 2 - Error : Will log all events that Fatal logs plus error events that might allow the application to continue running.
- 3 - Warn : Will log everything that Error logs plus other minor problems.
- 4 - Info : Will log everything that Warn logs plus informational messages that highlight the progress of the application at coarse-grained level.

- 5 - Info-With-Moab-Cmd-Debug : Will log everything that Info logs plus all the interaction with Moab Workload Manager EXCEPT the frequent (defaults to every 2 seconds) commands to determine if Moab Cluster Manager should refresh data from Moab Workload Manager.
- 6 - Debug : Will log everything that Info-With-Moab-Cmd-Debug level logs plus fine-grained informational events that are useful in debugging.
- 7 - Trace : The most verbose logging level. The only level that logs ALL interaction with Moab Workload Manager, including the frequent refresh checks ignored by other logging levels.

Users can configure extremely fine grained logging information by editing the <MCM_HOME>/conf/log4j.properties file. Using this file, one can set the log level on individual classes or packages within Moab Cluster Manager. It should be noted that configuring individual package or class loggers in the log4j.properties file other than the root logger or the logger for com.moab.api.XMLDebuggingTools may override the settings applied from this window.

Warning

Verbose log levels cause a small performance penalty. Because levels 5 and above log all the interaction with Moab Workload Manager, they can use substantially more memory when connected to larger systems.

8.4. Cluster Manager Preferences

These preferences control Cluster Manager-specific settings.

Refresh Rate

The more often Cluster Manager communicates with Workload Manager the more up-to-date the information; however, when Cluster Manager communicates with Workload Manager more often, Workload Manager operates at a slower pace.

- Fast - At this refresh rate Cluster Manager will update its information every minute.
- Medium - At this refresh rate, which is the default refresh rate, Cluster Manager will update its information every 10 minutes.
- Slow - At this refresh rate Cluster Manager will update its information every hour.

Advanced Settings

Check this box to enable more advanced options throughout Cluster Manager.

8.5. Plugin Manager

A plugin is a file which adds additional functionality to the Cluster Manager. They can be added and removed from the Plugin Manager. Click the folder icon next to 'Select Plugin File' and locate the plugin you wish to add. Hit the 'Load Plugin' button to load the plugin you have selected.

Chapter 9. License

This product was created by Cluster Resources, Inc. Copyright (C) 1999-2007 Cluster Resources, Inc., all rights reserved. Moab Workload Manager is a trademark of Cluster Resources, Inc.

This SOFTWARE is bound by an 'End User Open Source' LICENSE from Cluster Resources Inc. The conditions of the 'End User Open Source' LICENSE include, but are not limited to the conditions described below.

THE SOFTWARE IS PROVIDED AS IS, AND CLUSTER RESOURCES, INC. (CRI) AND ALL CONTRIBUTING PARTIES DISCLAIM ALL WARRANTIES RELATING TO THE SOFTWARE, WHETHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. NEITHER CRI NOR ANYONE INVOLVED IN THE CREATION, PRODUCTION, OR DELIVERY OF THE SOFTWARE SHALL BE LIABLE FOR ANY INDIRECT, CONSEQUENTIAL, OR INCIDENTAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE SOFTWARE EVEN IF CRI HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES OR CLAIMS. IN NO EVENT SHALL CRI'S LIABILITY FOR ANY DAMAGES EXCEED THE CONSIDERATION PAID FOR THE LICENSE TO USE THE SOFTWARE, REGARDLESS OF THE FORM OF CLAIM. THE PERSON OR ENTITY USING THE SOFTWARE BEARS ALL RISK AS TO THE QUALITY AND PERFORMANCE OF THE SOFTWARE.

By installing or using this SOFTWARE you are accepting a non-exclusive 'End User Open Source' LICENSE from Cluster Resources Inc. and are bound to abide by the following conditions:

- Inclusion of Notice and Disclaimer

  All copies of the SOFTWARE, whether or not for redistribution and whether or not in source code or in binary form must include a conspicuous and appropriate publication of the above copyright notice and disclaimer.

- Usage

  Source and/or binary forms of this SOFTWARE may be used by any 'End User' organization pursuant to the conditions of this and other associated LICENSES at no charge and for an unlimited period of time. An 'End User' organization is defined as an organization that is using this SOFTWARE on their own systems and is not commercially redistributing, modifying, supporting, or providing other services specific to this SOFTWARE to other organizations for profit.

- Modifications

  SOFTWARE may be freely modified by the 'End User' as necessary to meet the needs of the 'End User' LICENSEE'S system. 'End User' may solicit the services of Cluster Resources Inc. or

'Authorized Distribution and Services Partners' of Cluster Resources Inc. that have received express prior written authorization to redistribute, modify or provide services for SOFTWARE. Available services include but are not limited to technical support, training, consultation or optimization services. 'End User' may not solicit or receive this SOFTWARE or services associated to the use, customization, training, development, or support on this SOFTWARE from any organization that is not an 'Authorized Distribution and Services Partner' of Cluster Resources Inc. Any unauthorized partner that desires to become an 'Authorized Distribution and Services Partner' of Cluster Resources, Inc. may contact us at support@clusterresources.com. 'End User' organizations that desire services from Cluster Resources Inc., or an 'Authorized Distribution and Services Partner' may contact us using the same email listed above.

- Distribution

  'End User' organizations that are academic and government agencies may redistribute this SOFTWARE subject to the condition that the distribution contains conspicuous publication of the acknowledgement statement found within the LICENSE agreement distributed with this SOFTWARE.

  Organizations that are commercial and other for-profit organizations may not redistribute this code or derivations of this code in any form whatsoever, including parts of SOFTWARE incorporated into other software programs without express written permission from Cluster Resources, Inc.

  Redistribution of the SOFTWARE in any form whatsoever, including parts of the code that are incorporated into other software programs, must include a conspicuous and appropriate publication of the following acknowledgement:

  'This product was developed by Cluster Resources, Inc. Moab Scheduling System is a trademark of Cluster Resources, Inc.'

  Any such modification of the SOFTWARE must, when installed, display the above language, the copyright notice, and the warranty disclaimer.

  Each time the SOFTWARE is redistributed (or any work based on the SOFTWARE), the recipient must automatically receive this LICENSE, copyright notice, and the warranty disclaimer as described in this license agreement, which govern the ability to copy, distribute or modify the SOFTWARE subject to these terms and conditions, and has the choice of accepting or declining the LICENSE. As the LICENSEE, you shall automatically provide the recipient with a copy of this LICENSE. Further restrictions are not to be imposed on recipients of the SOFTWARE by the LICENSEE beyond those expressly described herein.

- Use of Modifications

  LICENSEES with a redistribution agreement that wish to distribute their modifications (including government and academic institutions) must first send a copy of the modifications along with a brief

explanation of why the modification was made and the resulting performance or functionality of the modifications to Cluster Resources, Inc. at support@clusterresources.com. Failure to send a copy of distributed modifications renders the LICENSE invalid, as well as any LICENSES granted to third parties subsequent to the incorporation of the modifications into SOFTWARE. Any such modification of the SOFTWARE must, when installed, display the LICENSE, the copyright notice, and the warranty disclaimer as described in the LICENSE agreement/s distributed with this SOFTWARE. Those without a LICENSE to redistribute may send modifications to Cluster Resources for evaluation and possible incorporation into SOFTWARE.

Copyright owners of modifications to SOFTWARE hereby grant Cluster Resources, Inc. a non-exclusive, royalty-free, worldwide, irrevocable right and LICENSE to install, use, distribute, sublicense, and prepare derivative works of said modifications. Only organizations receiving an express prior written exclusion to this condition are exempted from providing these non-exclusive rights to Cluster Resources, Inc.

- Communications about and Endorsement of SOFTWARE and Products/Software Derived from the SOFTWARE

  The name 'Moab Scheduling System' or 'Moab Scheduler' or any of its variants must not otherwise be used to endorse or to promote products derived from the SOFTWARE without prior written permission from CRI.

  Products derived from or incorporating the SOFTWARE in whole or in part shall not contain as part of the product's name any form of the terms 'Cluster Resources, Inc.', 'CRI', 'Moab', 'Moab Scheduling System', 'Moab Scheduler', or 'Supercluster Development Group' unless prior written permission has been received from Cluster Resources, Inc.

  All advertising materials for products that use or incorporate features of the SOFTWARE must display the following acknowledgement: 'This product includes software developed by Cluster Resources, Inc. for use in the Moab Scheduling System.'

- Acceptance of this LICENSE

  It is not required that you accept this LICENSE; however, if you do not accept the terms of this LICENSE, you are prohibited by law from installing, using, modifying or distributing the SOFTWARE or any of its derivative works. Therefore, by installing, using, modifying or distributing the SOFTWARE (or any of its derivative works), you have agreed to this LICENSE and have accepted all its terms and conditions.

  If any portion of this LICENSE is held invalid or unenforceable under any particular circumstance, the balance of the LICENSE will continue to apply.